A Provably Consistent Method for Imposing Sparsity in Feature-based SLAM Information Filters

Matthew Walter[1], Ryan Eustice[2], and John Leonard[1]

[1] Computer Science and Artificial Intelligence Laboratory
Massachusetts Institute of Technology
Cambridge MA, USA {mwalter,jleonard}@mit.edu
[2] Department of Applied Ocean Physics and Engineering
Woods Hole Oceanographic Institution
Woods Hole MA, USA ryan@whoi.edu *

Summary. An open problem in Simultaneous Localization and Mapping (SLAM) is the development of algorithms which scale with the size of the environment. A few promising methods exploit the key insight that representing the posterior in the canonical form parameterized by a sparse information matrix provides significant advantages regarding computational efficiency and storage requirements. Because the information matrix is naturally dense in the case of feature-based SLAM, additional steps are necessary to achieve sparsity. The delicate issue then becomes one of performing this sparsification in a manner which is consistent with the original distribution.

In this paper, we present a SLAM algorithm based in the information form in which sparseness is preserved while maintaining consistency. We describe an intuitive approach to controlling the population of the information matrix by essentially ignoring a small fraction of proprioceptive measurements whereby we track a modified version of the posterior. In this manner, the Exactly Sparse Extended Information Filter (ESEIF) performs exact inference, employing a model which is conservative relative to the standard distribution. We demonstrate our algorithm both in simulation as well as on two nonlinear datasets, comparing it against the standard EKF as well as the Sparse Extended Information Filter (SEIF) by Thrun *et al.* The results convincingly show that our method yields conservative estimates for the robot pose and map which are nearly identical to those of the EKF in comparison to the SEIF formulation which results in overconfident error bounds.

* R. Eustice is now with the Department of Mechanical Engineering at The Johns Hopkins University, Baltimore, MD USA, E-mail: rme@jhu.edu.

1 Introduction

A skill which plays an integral role in achieving robot autonomy is the ability to operate in *a priori* unknown environments. Viewed as a coupled problem of simultaneously performing localization and mapping, SLAM is further complicated by the stochastic nature of vehicle motion and observations. Most effective SLAM algorithms address these issues by posing the problem in a probabilistic framework with the goal then being the estimation of the joint distribution over the map and vehicle pose.

Beginning with the seminal work of Smith *et al.* [14], the Extended Kalman Filter (EKF) SLAM formulation has proven to be particularly popular. In large part, this is due to its relative simplicity, requiring that one only maintain the first two moments of the distribution to account for the coupling between the robot and map. From knowledge of the correlation, the EKF is able to exploit feature observation data to update the pose and map estimates. At the same time, this capability comes at the cost of complexity which is quadratic in the number of state elements. As a result, SLAM algorithms relying upon an EKF have traditionally been limited to relatively small environments.

Representing the joint Gaussian distribution in the dual canonical form, recent work has given rise to algorithms capable of scaling with the environment. Pivotal insights by Thrun *et al.* [15] and Frese *et al.* [8] have revealed that, in the context of SLAM, many of the off-diagonal elements in the inverse covariance (information) matrix are inherently near zero. Considering the graphical model represented by the information matrix [12], the implication is that a majority of the links in the Markov network are relatively weak. By essentially breaking these weak links, Frese [7] and Paskin [12] are able to approximate the graphical model by a sparse tree structure which provides for scalable SLAM algorithms. Alternatively, the Sparse Extended Information Filter (SEIF) by Thrun *et al.* [15] relies upon a version of the Extended Information Filter, the dual to the EKF. In the case where the information matrix is sparse, the authors demonstrate that state estimation can be performed in near-constant time. While a majority of the links in the information matrix are weak, though, they are nonetheless nonzero. SEIFs then employ a strategy by which the posterior is approximated with an information matrix having the desired sparse structure. The algorithm has efficiently been applied to large, real-world datasets with *a priori* unknown data association [10].

Together with the intuitive characteristics of the canonical representation noted in [15], the success of SEIFs has brought a lot of attention to the information filter formulation of the SLAM problem. The one issue which has, up to now, largely gone unnoticed is the implication of approximating the posterior to achieve the necessary sparseness. A close look at the sparsification strategy reveals that the resulting posterior is prone to overconfidence. In [5], the authors show that, while the state estimates are only slightly overconfident when expressed in a local reference frame, they suffer from an exaggerated global inconsistency. The paper presents a modified sparsification rule which

yields a posterior which is both locally and globally consistent relative to the full Kalman solution but is no longer computationally tractable.

Our objective in this paper is to present an information-based formulation to the SLAM problem which achieves exact sparseness while being computationally efficient. Rather than relying upon an approximation to remove links from the information matrix, the algorithm adopts a new strategy which actively controls the population of the matrix by relocalizing the robot within the map. The filter then maintains an estimate of the state which is both globally *and* locally conservative relative to the full Kalman solution. We demonstrate the algorithm alongside the EKF and SEIF on a linear Gaussian simulation as well as two real-world experiments, including a benchmark nonlinear dataset. The results reveal that while the SEIF is globally inconsistent, our algorithm yields estimates nearly identical to those of the EKF which are globally and locally conservative.

2 Information Filter

2.1 Canonical Form

Let $\boldsymbol{\xi}_t$ be a random vector having a Gaussian probability density, $\boldsymbol{\xi}_t \sim \mathcal{N}(\boldsymbol{\xi}_t; \boldsymbol{\mu}_t, \Sigma_t)$ described completely by its mean, $\boldsymbol{\mu}_t$, and covariance matrix, Σ_t. An expansion of the exponential term defining the multivariate normal distribution, $p(\boldsymbol{\xi}_t) \propto \exp\left\{-\frac{1}{2}(\boldsymbol{\xi}_t - \boldsymbol{\mu}_t)^\top \Sigma_t^{-1}(\boldsymbol{\xi}_t - \boldsymbol{\mu}_t)\right\}$, yields an equivalent representation for the probability density function, $\mathcal{N}^{-1}(\boldsymbol{\xi}_t; \boldsymbol{\eta}_t, \Lambda_t)$, parameterized by the information vector and information matrix, $\boldsymbol{\eta}_t$ and Λ_t, respectively.

$$\Lambda_t = \Sigma_t^{-1} \quad \boldsymbol{\eta}_t = \Sigma_t^{-1}\boldsymbol{\mu}_t \tag{1}$$

The canonical representation for the multivariate Gaussian is the dual of the standard form in the sense of the fundamental processes of marginalization and conditioning, as exemplified in Table 1. While marginalization is *hard* in the information form, requiring a matrix inversion, it is *easy* in the covariance form. The opposite is true in regards to the conditioning operation. Further details regarding this duality in the context of filtering can be found in [11].

One quality of the canonical form is its relationship with Gaussian Markov random fields in which nodes in the graph represent individual state variables and edge structure describes their conditional independence relationships. The information matrix effectively serves as an adjacency matrix for the graph [12], with the strength of constraints between pairs of variables proportional to the corresponding elements of the matrix. Off-diagonal components which are zero then denote the absence of links in the Markov network. Thus, the information matrix has the particular advantage of explicitly representing the conditional independence of state variables.

Table 1. Summary of Marginalization and Conditioning Operations on a Gaussian Distribution Expressed in Covariance and Information Form

$$p(\boldsymbol{\alpha},\boldsymbol{\beta}) = \mathcal{N}\left(\begin{bmatrix}\boldsymbol{\mu}_\alpha \\ \boldsymbol{\mu}_\beta\end{bmatrix}, \begin{bmatrix}\Sigma_{\alpha\alpha} & \Sigma_{\alpha\beta} \\ \Sigma_{\beta\alpha} & \Sigma_{\beta\beta}\end{bmatrix}\right) = \mathcal{N}^{-1}\left(\begin{bmatrix}\boldsymbol{\eta}_\alpha \\ \boldsymbol{\eta}_\beta\end{bmatrix}, \begin{bmatrix}\Lambda_{\alpha\alpha} & \Lambda_{\alpha\beta} \\ \Lambda_{\beta\alpha} & \Lambda_{\beta\beta}\end{bmatrix}\right)$$

	Marginalization $p(\boldsymbol{\alpha}) = \int p(\boldsymbol{\alpha},\boldsymbol{\beta})\, d\boldsymbol{\beta}$	Conditioning $p(\boldsymbol{\alpha} \mid \boldsymbol{\beta}) = p(\boldsymbol{\alpha},\boldsymbol{\beta}) / p(\boldsymbol{\beta})$
Covariance Form	$\boldsymbol{\mu} = \boldsymbol{\mu}_\alpha$ $\Sigma = \Sigma_{\alpha\alpha}$	$\boldsymbol{\mu}' = \boldsymbol{\mu}_\alpha + \Sigma_{\alpha\beta}\Sigma_{\beta\beta}^{-1}(\boldsymbol{\beta} - \boldsymbol{\mu}_\beta)$ $\Sigma' = \Sigma_{\alpha\alpha} - \Sigma_{\alpha\beta}\Sigma_{\beta\beta}^{-1}\Sigma_{\beta\alpha}$
Information Form	$\boldsymbol{\eta} = \boldsymbol{\eta}_\alpha - \Lambda_{\alpha\beta}\Lambda_{\beta\beta}^{-1}\boldsymbol{\eta}_\beta$ $\Lambda = \Lambda_{\alpha\alpha} - \Lambda_{\alpha\beta}\Lambda_{\beta\beta}^{-1}\Lambda_{\beta\alpha}$	$\boldsymbol{\eta}' = \boldsymbol{\eta}_\alpha - \Lambda_{\alpha\beta}\boldsymbol{\beta}$ $\Lambda' = \Lambda_{\alpha\alpha}$

2.2 Feature-based SLAM

The goal of any SLAM algorithm is to concurrently perform navigation and map-building in the presence of uncertainty in vehicle motion and environmental observations. With feature-based SLAM formulations, the map is described as a collection of stationary primitives, e.g. lines, points, etc. The robot pose, $\mathbf{x}_t$, together with the set of map elements, $\mathbf{M} = \{\mathbf{m}_1, \mathbf{m}_2, \ldots, \mathbf{m}_n\}$, are represented together by the state vector, $\boldsymbol{\xi}_t = [\mathbf{x}_t^\top \ \mathbf{M}^\top]^\top$. The coupling between the pose and map is addressed by considering the joint probability distribution for the state. Adopting a Bayesian framework, a model of the joint posterior is tracked as it evolves as a result of the uncertainty in vehicle motion and measurement data. Typical SLAM implementations make the assumption that this uncertainty is a result of independent white Gaussian noise which corrupts the motion and measurement models. One can then show that the posterior obeys a Gaussian distribution.

$$p\left(\boldsymbol{\xi}_t | \mathbf{z}^t, \mathbf{u}^t\right) = \mathcal{N}\left(\boldsymbol{\xi}_t; \boldsymbol{\mu}_t, \Sigma_t\right) = \mathcal{N}^{-1}\left(\boldsymbol{\xi}_t; \boldsymbol{\eta}_t, \Lambda_t\right) \tag{2}$$

The belief function is traditionally represented in the standard form which can be tracked relatively easily with the EKF. Modifying the posterior to reflect the effect of vehicle motion is a constant-time process as it involves a combined process of state augmentation and marginalization, both of which are easily performed in the covariance form. On the other hand, it is well known that incorporating new measurement data requires a conditioning step which is quadratic in the size of the state. Furthermore, maintaining the correlation among state estimates leads to a dense covariance matrix which must be stored. For small scale environments, these problems are surmountable, but as the map size becomes increasingly large, implementing a full EKF quickly becomes intractable.

Alternatively, employing the canonical representation of the posterior, the filtering process reflects the duality between the two forms. Performing measurement updates (conditioning) is constant-time, while the marginalization component of the projection step, in general, is quadratic in the state dimension, at best. Furthermore, recovering the estimate of the mean requires the $\mathcal{O}(n^3)$ inversion of the information matrix per (1). As a result of these limitations, the information filter has had relatively limited use in SLAM.

Recently, Thrun *et al.* [15] and Frese *et al.* [8] have made the pivotal observation that, when normalized, the information matrix tends to be nearly sparse. The matrix is dominated by a small percentage of terms which are significantly larger than the remaining elements. In general, the links between the robot and the map are stronger for nearby, recently observed features while the constraints are weak for distant features. The same is true for inter-landmark terms which tend to decay exponentially with the distance traversed by the robot between observations [6]. Referring to the graphical interpretation of the information matrix, these weak links then imply that, given relatively few features, the robot is nearly conditionally independent of much of the map.

Though many of the terms in the normalized information matrix are very small, the SLAM process naturally leads to the full population of the matrix. To get a better understanding of why this is, consider a simple example in which the map consists of five features. Suppose that the off-diagonal terms in the information matrix corresponding to the robot, $\mathbf{x}_t$, are non-zero for four of the features and that the remaining landmark, $\mathbf{m}_4$, has shared information with another feature. These links between the robot and the map are created when features are observed. The graphical model along with the information matrix are illustrated in the left-hand side of Figure 1(a). The time projection step can be viewed as an initial augmentation of the state with the new robot pose, $\mathbf{x}_{t+1}$, which, evolving by a Markov process, is linked only to the previous pose as indicated in the middle figure. At this point, the information matrix remains sparse. Subsequently marginalizing out $\mathbf{x}_t$, though, creates links between all states which share constraints with the previous pose. The result is a fully connected subset of nodes and, correspondingly, a population of the information matrix. The only remaining zero entries correspond to the lone feature, $\mathbf{m}_4$, which will become linked to the robot upon the next observation. The time projection step will then lead to a fully connected graph and, correspondingly, a dense information matrix.

Hence, with *online* SLAM implementations in which only the current pose of the robot is estimated, the marginalization of the previous pose in the projection step naturally results in a dense information matrix. Alternatively, by retaining an entire trajectory history, exact sparsity can be maintained [3] at the cost of storage requirements which become significant for large datasets.

Returning to the example pictoralized in Figure 1(a), note that while the time projection step populates the information matrix, the strength of the off-diagonal links decays with time. This behavior is the reason why a majority of the elements in the normalized matrix are very small. The authors show

in [15, 8, 12] that if the posterior can be represented by an exactly sparse approximation, it is possible to achieve significant gains when it comes to both storage and time requirements. In particular, a bound on the number of links between the robot and the map allows for near constant-time performance of the time projection step and also controls the fill-in of the information matrix resulting from marginalization.

3 Exactly Sparse Extended Information Filters

Map elements having shared information with the robot are said to be *active*. In feature-based SLAM implementations, a feature becomes active when it is first observed. With time, the strength of the link with the robot decays and is strengthened only upon being re-observed. Thus, while the off-diagonal terms may become arbitrarily small, they will never become zero. In order for a landmark to become *passive* (i.e. no shared information), the link with the robot must explicitly be broken.

3.1 Problem Formulation

In describing the desired sparsity of the information matrix, we adopt the two measures utilized by Thrun *et al.* [15]. Denote the maximum allowable number of active features as Γ_a and the number of inter-landmark links in the matrix by Γ_p. Let us then partition the map into two sets of features, $\mathbf{M} = \{\mathbf{m}^+, \mathbf{m}^-\}$, where $\mathbf{m}^+$ represents the active features for which the off-diagonal terms for the robot pose are non-zero, and $\mathbf{m}^-$ denotes the passive landmarks, having no direct constraint to the vehicle.

Controlling the sparsity of the information matrix is, in large part, a direct consequence of maintaining the Γ_a bound. By regulating the number of active features, it is possible to limit the population of the matrix. Consider, for example, the situation depicted in the left-hand side of Figure 1(b) in which four of the five features are active. At this point, if $\mathbf{x}_t$ were marginalized out, the four active features in $\mathbf{m}^+$ would become fully connected, potentially violating the Γ_p bound. Instead, if one of landmarks, $\mathbf{m}_1$, were first made passive, the number of non-zero elements created as a result of marginalization could be controlled. Thus, enforcing the desired sparsity pattern corresponds to maintaining a bound on the number of active features. Since features do not naturally become passive, a sparsification routine which deliberately breaks the links is necessary.

3.2 SEIF Sparsification

The SEIF breaks a link between the robot and a feature by approximating the posterior with a distribution in which the robot is conditionally independent of

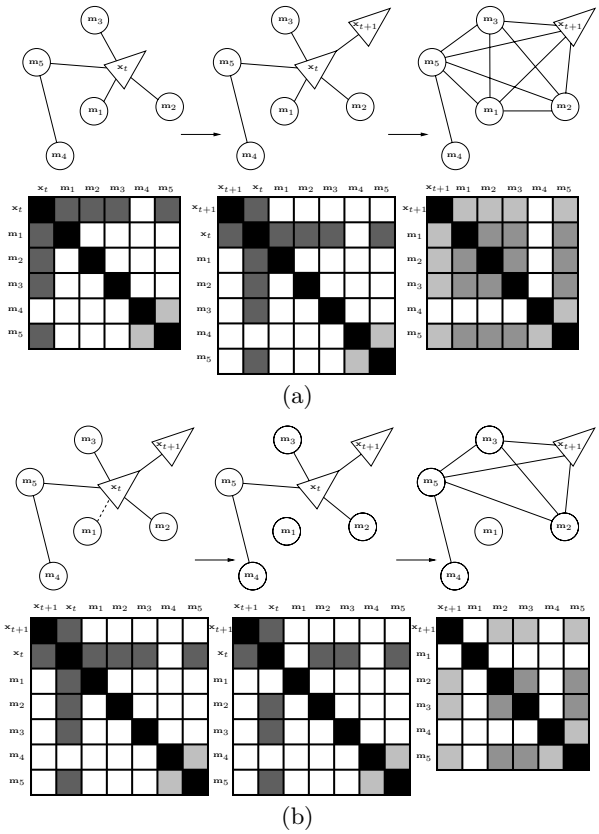

Fig. 1. A graphical explanation of SEIF's methodology for controlling sparsity in the information matrix. (a) A sequence of illustrations depicting the evolution of the Markov network and corresponding information matrix resulting from time projection when viewed as a two-step process of state augmentation followed by marginalization. Darker shades imply larger magnitudes with white indicating zero values. From left to right we have: (1) the robot $\mathbf{x}_t$ connected to four active features, $\mathbf{m}_{1:3}$ and $\mathbf{m}_5$; (2) state augmentation of the time-propagated robot pose $\mathbf{x}_{t+1}$; (3) marginalized distribution where the old pose, $\mathbf{x}_t$, has been eliminated. (b) A sequence of illustrations highlighting the concept behind sparsification. If feature $\mathbf{m}_1$ can first be made passive by eliminating its link to the old pose, $\mathbf{x}_t$, then marginalization over $\mathbf{x}_t$ will not link it to the other active features. This implies that we can control fill-in of the information matrix by bounding the number of currently active features.

the landmark. The map is broken into three disjoint sets, $\mathbf{M} = \{\mathbf{m}^0, \mathbf{m}^+, \mathbf{m}^-\}$, where $\mathbf{m}^-$ refers to the passive landmarks which will remain passive and, in a slight abuse of notation, $\mathbf{m}^+$ is the set of active features which will remain active, and $\mathbf{m}^0$ are the active features which will be made passive. The sparsification routine proceeds from a decomposition of the posterior

$$\begin{aligned} p(\mathbf{x}_t, \mathbf{m}^0, \mathbf{m}^+, \mathbf{m}^-) &= p(\mathbf{x}_t \mid \mathbf{m}^0, \mathbf{m}^+, \mathbf{m}^-)\, p(\mathbf{m}^0, \mathbf{m}^+, \mathbf{m}^-) \\ &= p(\mathbf{x}_t \mid \mathbf{m}^0, \mathbf{m}^+, \mathbf{m}^- = S_{\mathbf{m}^-}^\top \boldsymbol{\mu}_t)\, p(\mathbf{m}^0, \mathbf{m}^+, \mathbf{m}^-) \end{aligned} \tag{3}$$

where setting the passive elements to their mean, $S_{\mathbf{m}^-}^\top \boldsymbol{\mu}_t$, in the last line is valid due to their conditional independence with respect to the robot. SEIFs then deactivate the landmarks by replacing (3) with an approximation to the posterior which drops the dependence upon $\mathbf{m}^0$:

$$\tilde{p}_{\text{SEIF}}(\mathbf{x}_t, \mathbf{m}^0, \mathbf{m}^+, \mathbf{m}^-) = p(\mathbf{x}_t \mid \mathbf{m}^+, \mathbf{m}^- = S_{\mathbf{m}^-}^\top \boldsymbol{\mu}_t)\, p(\mathbf{m}^0, \mathbf{m}^+, \mathbf{m}^-) \tag{4}$$

While the decomposition in (3) is theoretically sound, it is no longer valid to condition on a particular value for the passive features while simultaneously ignoring the dependence upon $\mathbf{m}^0$. Given only a subset of the active features, the robot pose is no longer conditionally independent of the passive map.

By enforcing the conditional independence between the robot and the deactivated features, SEIFs rely upon approximate inference on an approximate posterior and, as a result, are prone to inconsistency [5]. In particular, the authors show that sparsifying in this manner leads to a *global* map which is significantly overconfident while the *local* relationships are preserved.

3.3 ESEIF Sparsification

Rather than deliberately breaking constraints with the robot to maintain a bound on the number of active features, ESEIFs take the approach of essentially controlling the initial formation of links. As soon as a feature is first observed, it is linked to the current robot pose. As noted earlier, the strength of this constraint will decay with time but never truly disappear, leading to a growing number of links between the robot and the map.

Noting the nature of this link formation, ESEIFs control the number of active features by deliberately marginalizing out the robot pose. The vehicle is relocated within the map using observations of a few known landmarks. The new pose is then conditionally independent of the rest of the map, and the robot is linked only to the features used for relocalization.

For a more detailed description of the ESEIF sparsification strategy, we consider a situation which would give rise to the representation in Figure 1 which consists of both active and passive features. Suppose that the robot makes four observations, $\mathcal{Z}_t = \{\mathbf{z}_1, \mathbf{z}_2, \mathbf{z}_3, \mathbf{z}_5\}$, three being of active features and one of a passive feature:

$$\begin{aligned} \mathbf{z}_2 &= \mathbf{h}(\mathbf{x}_v, \mathbf{m}_2),\ \mathbf{m}_2 \in \mathbf{m}^+ & \mathbf{z}_5 &= \mathbf{h}(\mathbf{x}_v, \mathbf{m}_5),\ \mathbf{m}_5 \in \mathbf{m}^+ \\ \mathbf{z}_1 &= \mathbf{h}(\mathbf{x}_v, \mathbf{m}_1),\ \mathbf{m}_1 \in \mathbf{m}^0 & \mathbf{z}_3 &= \mathbf{h}(\mathbf{x}_v, \mathbf{m}_3),\ \mathbf{m}_3 \in \mathbf{m}^- \end{aligned}$$

Updating the posterior based upon all four measurements would result in the strengthening of the off-diagonal entries in the matrix pairing the robot with the three observed active features. Additionally, a link would be created

with the currently passive map element, $\mathbf{m}_3$, leading to the graph structure depicted in the left-hand side of Figure 1(a). In the case where this would lead to a violation of the Γ_a bound, one strategy would be to disregard the observation of the passive feature entirely. With ESEIFs, though, it is possible to incorporate all measurement data while maintaining the desired sparsity pattern.

In the ESEIF sparsification step, the measurement data is partitioned into two sets, $\mathbf{z}_\alpha$ and $\mathbf{z}_\beta$, where the first set is used for updating the filter and the second is reserved for performing relocalization. Of the four measurements available in our example, group that of the passive feature together with one active measurement as $\mathbf{z}_\alpha = \{\mathbf{z}_1, \mathbf{z}_3\}$, leaving $\mathbf{z}_\beta = \{\mathbf{z}_2, \mathbf{z}_5\}$. To sparsify, we first apply the update step followed by the combined process of marginalization and relocalization.

Posterior Update

A Bayesian update is performed on the joint posterior, $p(\boldsymbol{\xi}_t \mid \mathbf{z}^{t-1}, \mathbf{u}^t) = \mathcal{N}^{-1}\big(\boldsymbol{\xi}_t; \boldsymbol{\eta}_t, \Lambda_t\big)$ based upon the $\mathbf{z}_\alpha$ measurements:

$$p(\boldsymbol{\xi}_t \mid \mathbf{z}^{t-1}, \mathbf{u}^t) \xrightarrow{\mathbf{z}_\alpha = \{\mathbf{z}_1, \mathbf{z}_3\}} p_1\big(\boldsymbol{\xi}_t \mid \{\mathbf{z}^{t-1}, \mathbf{z}_\alpha\}, \mathbf{u}^t\big)$$

where $p_1\big(\boldsymbol{\xi}_t \mid \{\mathbf{z}^{t-1}, \mathbf{z}_\alpha\}, \mathbf{u}^t\big) = \mathcal{N}^{-1}\big(\boldsymbol{\xi}_t; \bar{\boldsymbol{\eta}}, \bar{\Lambda}\big)$ follows from the standard update process for the information filter. Note that we can perform this step in constant-time with, in the nonlinear case, access to the mean estimate for the robot as well as $\mathbf{m}_1$ and $\mathbf{m}_3$. The information matrix, $\bar{\Lambda}_t$, is modified as depicted in Figure 2 with the strengthening of the constraints between the vehicle and the active feature, $\mathbf{m}_1$ and importantly, the creation of shared information with the previously passive feature, $\mathbf{m}_3$.

Marginalization and Relocalization

The addition of a new constraint between the robot and a map element results in a violation of, Γ_a, the bound on the number of active features. The ESEIF sparsification routine then proceeds by first marginalizing out the vehicle pose

$$\begin{aligned} p_2\big(\mathbf{M}_t \mid \{\mathbf{z}^{t-1}, \mathbf{z}_\alpha\}, \mathbf{u}^t\big) &= \int_{\mathbf{x}_t} p_1\big(\boldsymbol{\xi}_t \mid \{\mathbf{z}^{t-1}, \mathbf{z}_\alpha\}, \mathbf{u}^t\big) d\mathbf{x}_t \\ &= \mathcal{N}^{-1}\big(\mathbf{M}_t; \check{\boldsymbol{\eta}}_t, \check{\Lambda}_t\big) \end{aligned}$$

Following the representation of the marginalization process presented in Table 1, the canonical parameterization of the marginal is calculated as

$$p_2\big(\mathbf{M}_t \mid \{\mathbf{z}^{t-1}, \mathbf{z}_\alpha\}, \mathbf{u}^t\big) \;=\; \mathcal{N}^{-1}\big(\mathbf{M}_t; \check{\boldsymbol{\eta}}_t, \check{\Lambda}_t\big)$$

$$\begin{aligned}\check{\Lambda}_t = {} & \mathrm{S}^\top_{\mathbf{m}^0,\mathbf{m}^+,\mathbf{m}^-} \bar{\Lambda}_t \mathrm{S}_{\mathbf{m}^0,\mathbf{m}^+,\mathbf{m}^-} \\ & - \mathrm{S}^\top_{\mathbf{m}^0,\mathbf{m}^+,\mathbf{m}^-} \bar{\Lambda}_t \mathrm{S}_{\mathbf{x}_t} \big(\mathrm{S}^\top_{\mathbf{x}_t} \bar{\Lambda}_t \mathrm{S}_{\mathbf{x}_t}\big)^{-1} \mathrm{S}^\top_{\mathbf{x}_t} \bar{\Lambda}_t \mathrm{S}_{\mathbf{m}^0,\mathbf{m}^+,\mathbf{m}^-}\end{aligned} \tag{5a}$$

$$\check{\boldsymbol{\eta}}_t = \mathrm{S}^\top_{\mathbf{m}^0,\mathbf{m}^+,\mathbf{m}^-} \bar{\boldsymbol{\eta}}_t - \mathrm{S}^\top_{\mathbf{m}^0,\mathbf{m}^+,\mathbf{m}^-} \bar{\Lambda}_t \mathrm{S}_{\mathbf{x}_t} \big(\mathrm{S}^\top_{\mathbf{x}_t} \bar{\Lambda}_t \mathrm{S}_{\mathbf{x}_t}\big)^{-1} \mathrm{S}^\top_{\mathbf{x}_t} \bar{\boldsymbol{\eta}}_t \tag{5b}$$

where $\mathrm{S}_{\mathbf{m}^0,\mathbf{m}^+,\mathbf{m}^-}$ and $\mathrm{S}_{\mathbf{x}_t}$ are projection matrices mapping the state space to the $\{\mathbf{m}^0, \mathbf{m}^+, \mathbf{m}^-\}$ and $\mathbf{x}_t$ subspaces, respectively.

The inverse term involves the block diagonal of the information matrix corresponding to the vehicle pose, $\mathrm{S}^T_{\mathbf{x}_t} \bar{\Lambda}_t \mathrm{S}_{\mathbf{x}_t}$, which is of fixed size. Meanwhile, the $\mathrm{S}^T_{\mathbf{m}^0,\mathbf{m}^+,\mathbf{m}^-} \bar{\Lambda}_t \mathrm{S}_{\mathbf{x}_t}$ matrix corresponds to the shared information between the map and the vehicle pose and, taken as an outer product over the vehicle sub-block, yields a matrix having nonzero values only for the active feature indices. It is a result of this term that marginalization establishes the connectivity among the active features shown in the right-hand side of Figure 2. The computational complexity of this matrix outer product is limited by the Γ_a bound and the order of the matrix inversion is fixed. Thus, the marginalization can be performed in constant-time.

We complete sparsification in ESEIFs by relocalizing the vehicle within the map using the remaining $\mathbf{z}_\beta$ measurements. The new pose estimate is, in general, given by a nonlinear function of measurement data and corresponding feature estimates of the form in (6a) where $\mathbf{w}_t \sim \mathcal{N}\big(\mathbf{w}_t; \mathbf{0}, \mathrm{R}\big)$ is white Gaussian noise. Equation (6b) corresponds to the linearization about the mean of the marginal distribution, $\mathcal{N}^{-1}\big(\mathbf{M}_t; \check{\boldsymbol{\eta}}_t, \Lambda\big)$ in (5). The Jacobian, G, is sparse as the only non-zero columns are those corresponding to the map elements used for relocalization. Subsequently, only the mean estimates for these features are necessary for the linearization.

$$\mathbf{x}_t = \mathbf{g}\big(\mathbf{m}_\beta, \mathbf{z}_\beta\big) + \mathbf{w}_t \tag{6a}$$

$$\approx \mathbf{g}\big(\check{\boldsymbol{\mu}}_{m_\beta}, \mathbf{z}_\beta\big) + \mathrm{G}\big(\mathbf{m} - \check{\boldsymbol{\mu}}_t\big) + \mathbf{w}_t \tag{6b}$$

Augmenting the map distribution (5) with the new pose estimate yields a state which can be shown to have the following canonical parameterization:

$$p_{\mathrm{ESEIF}}(\boldsymbol{\xi}_t \mid \mathbf{z}^t, \mathbf{u}^t) = \mathcal{N}^{-1}\big(\boldsymbol{\xi}_t; \breve{\boldsymbol{\eta}}_t, \breve{\boldsymbol{\Lambda}}_t\big)$$

$$\breve{\boldsymbol{\eta}}_t = \begin{bmatrix} \mathrm{R}^{-1}\big(\mathbf{g}(\check{\boldsymbol{\mu}}_{m_\beta}, \mathbf{z}_\beta) - \mathrm{G}\check{\boldsymbol{\mu}}_t\big) \\ \check{\boldsymbol{\eta}}_t - \mathrm{G}^\top \mathrm{R}^{-1}\big(\mathbf{g}(\check{\boldsymbol{\mu}}_{m_\beta}, \mathbf{z}_\beta) - \mathrm{G}\check{\boldsymbol{\mu}}_t\big) \end{bmatrix} \tag{7a}$$

$$\breve{\Lambda}_t = \begin{bmatrix} \mathrm{R}^{-1} & -\mathrm{R}^{-1}\mathrm{G} \\ -\mathrm{G}^\top \mathrm{R}^{-1} & \big(\check{\Lambda}_t + \mathrm{G}^\top \mathrm{R}^{-1}\mathrm{G}\big) \end{bmatrix} \tag{7b}$$

Due to the sparsity of G, most terms in $-\mathrm{R}^{-1}\mathrm{G}$ of the information matrix in (7b) that link the robot to the map are zero, except for those corresponding to the landmarks used for relocalization. The new instantiation for the

robot pose is then conditionally independent of the rest of the map. As a result, ESEIF sparsification leads to the joint posterior having the desired factorization:

$$p_{\text{ESEIF}}(\boldsymbol{\xi}_t \mid \mathbf{z}^t, \mathbf{u}^t) = p(\mathbf{x}_t \mid \mathbf{m}_\beta, \mathbf{z}_\beta)\, p_2(\mathbf{M}_t \mid \{\mathbf{z}^{t-1}, \mathbf{z}_\alpha\}, \mathbf{u}^t)$$

As reflected by the resulting information matrix depicted in Figure 2, the active features are then limited to those used for relocalization.

In this manner, ESEIFs control the size of the active map and, in turn, the sparseness of the information matrix. Like the full EKF, the ESEIF performs exact inference on an approximate model, albeit on a different posterior. When we first marginalize out (kidnap) and subsequently relocalize the robot, we are performing the dual of kidnapping and relocation for the standard EKF. Essentially, we are ignoring the odometry data which links the current and previous poses. Hence, whereas the full EKF tracks the Gaussian approximation to the posterior, $p(\boldsymbol{\xi}_t \mid \mathbf{Z}^t)$, ESEIFs and the relocated EKF maintain the Gaussian model of an alternate distribution, $p(\boldsymbol{\xi}_t \mid \mathbf{Z}^{t^*})$. In this way, the ESEIF employs exact inference on an approximate model for which the information matrix is exactly sparse.

3.4 Recovering the Mean

A drawback of representing the posterior in the canonical form is that we no longer have access to the mean vector or covariance matrix. When the system equations are nonlinear, a subset of the mean is required to perform linearizations. Naively, we could recover the entire mean vector as $\boldsymbol{\mu}_t = \Lambda_t^{-1}\boldsymbol{\eta}_t$, though this operation is cubic in the dimension of the state and quickly becomes intractable. Instead, we can pose the problem in terms of solving a set of linear equations

$$\Lambda_t \boldsymbol{\mu}_t = \boldsymbol{\eta}_t \tag{8}$$

and take advantage of the sparseness of the information matrix. There are a number of techniques which iteratively solve such sparse, symmetric positive definite systems including conjugate gradient descent [13] and, more recently, the multilevel method proposed by [9]. Aside from loop closures, the mean vector evolves rather slowly in SLAM and, thus, the optimization can be performed over the course of multiple time steps. This then allows us to bound the number of iterations required per time step [2].

3.5 Data Association

Traditionally, the problem of data association is addressed by evaluating the likelihood of a measurement for different correspondence hypothesis. The distribution follows from marginalizing out all state elements except for the variables we are interested in, $\mathbf{x}_i$ and $\mathbf{x}_j$. From the duality indicated in Table 1,

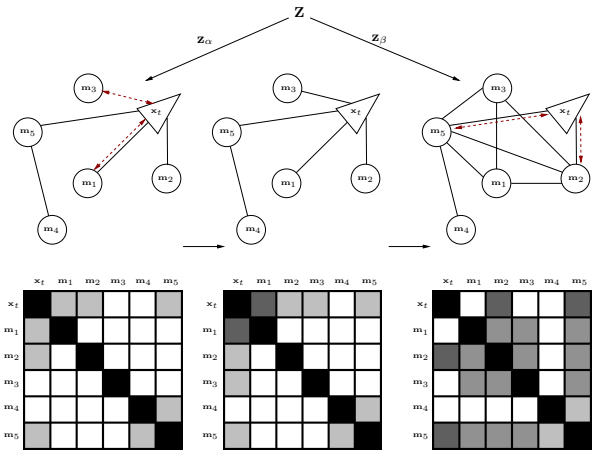

Fig. 2. Sparsification as performed by ESEIFs during the measurement update step. At time t, three of the mapped features are active, $\mathbf{m}^+ = \{\mathbf{m}_1, \mathbf{m}_2, \mathbf{m}_5\}$ and two are passive, $\mathbf{m}^- = \{\mathbf{m}_3, \mathbf{m}_4\}$ as indicated by shaded off-diagonal elements of the information matrix. The robot makes three observations of active features, $\mathbf{z}_1$, $\mathbf{z}_2$, and $\mathbf{z}_5$, and one of a passive feature, $\mathbf{z}_3$. The first step of the ESEIF sparsification algorithm, as shown in the left-most diagram, is to update the posterior based upon a subset of the measurements, $\mathbf{z}_\alpha = \{\mathbf{z}_1, \mathbf{z}_3\}$, resulting in a stronger constraint with $\mathbf{m}_1$ as well as the formation of a link with $\mathbf{m}_3$, as depicted in the middle figure. Sparsification then proceeds with the marginalization of the vehicle pose and subsequent relocation of the robot based upon the remaining measurements, $\mathbf{z}_\beta$. The implication on the information matrix is the connectivity of the initial set of active features and a desired restriction on the number of constraints with the vehicle pose.

this operation is easy in the standard form but difficult with the canonical parameterization where a large matrix inversion is necessary. Instead, Thrun *et al.* [15] first compute the conditional distribution for the Markov blanket for $\mathbf{x}_i$ and $\mathbf{x}_j$, $p(\mathbf{x}_i, \mathbf{x}_j, \mathbf{x}_k \mid \mathbf{x}_l)$, which involves simply extracting a sub-block of the information matrix. They then invert this matrix and take the covariance sub-block corresponding to $p(\mathbf{x}_i, \mathbf{x}_j \mid \mathbf{x}_l)$ which they use for data association. While the authors have had success using this conditional covariance, it can be shown to yield overconfident estimates for the likelihood [4].

Alternatively, Eustice *et al.* [4] propose a method which solves for conservative estimates for the marginal covariance. The technique stems from posing the relationship, $\Lambda_t \Sigma_t = \mathrm{I}$, as a sparse system of linear equations, $\Lambda_t \Sigma_{*i} = \mathbf{e}_i$, where Σ_{*i} and $\mathbf{e}_i$ denote the i^{th} columns of the covariance and identity matrices, respectively. To determine the robot pose covariance, the iterative algorithms previously presented for mean recovery can be used to solve the set of equations formed from the robot pose columns. Combining the estimate for robot pose covariance with a conservative estimate for the covariance of any map element gives rise to a joint covariance which is it-

self conservative. The joint covariance is then used to represent the marginal distribution for data association.

4 Results

To better understand the effectiveness of the two different sparse information filters, we compare the performance of ESEIFs and SEIFs to the standard EKF when applied to different forms of the SLAM problem. In the first case, we take a look at a controlled linear Gaussian simulation for which the KF, the optimal Bayesian estimator, is the "gold standard". We then follow with experiments using real-world nonlinear datasets including a benchmark outdoor data set widely popular in the SLAM community.

4.1 Linear Gaussian Simulation

To systematically analyze the two information-based filters, we first apply the three estimators in a controlled simulation. The environment consists of a set of point features, uniformly distributed to achieve a desired density of 0.10 features per unit area. The vehicle moves translationally according to a linear, constant velocity motion model and, at any time step, is able to observe the relative position to a limited number of neighboring features. Both the vehicle motion as well as the measurements are corrupted by additive white Gaussian noise.

As a basis for comparison, we apply the Kalman Filter, the optimal estimator for linear Gaussian problems. The ESEIF and SEIF are implemented with a limit of $\Gamma_a = 10$ active features. When sparsifying the ESEIF, we reserve as many of the observations for relocalizing the robot as possible, to the extent that we do not violate the Γ_a bound.

In the LG case, sparse information filters have already been shown to be computationally efficient [15]. Instead, we are interested in evaluating the effect that the different sparsification strategies have on the estimation accuracy. To that end, we perform a series of Monte Carlo simulations, using the normalized estimation error squared (NEES) [1] to measure filter consistency using a pair of metrics. As one measure, we use the Euclidean distance between the state estimates and the ground truth which corresponds to the *global* error. To get a *local/relative* measure of error, we first reference the robot and map positions relative to the first observed feature, $\mathbf{x}_m$ using the standard compounding operation, $\mathbf{x}_{mi} = \ominus\mathbf{x}_m \oplus \mathbf{x}_i$. We then compute the second error metric as the distance to the root-shifted representation of the ground truth. We plot the global normalized errors for the estimated vehicle position as well as for one of the map elements in Figures 3(a) and 3(b), respectively. Comparing these errors with the $97.5^{\%}$ chi-square upper bound indicated by the horizontal line, we see that the ESEIF yields consistent position estimates with errors similar to those of the KF. The normalized errors attributed to

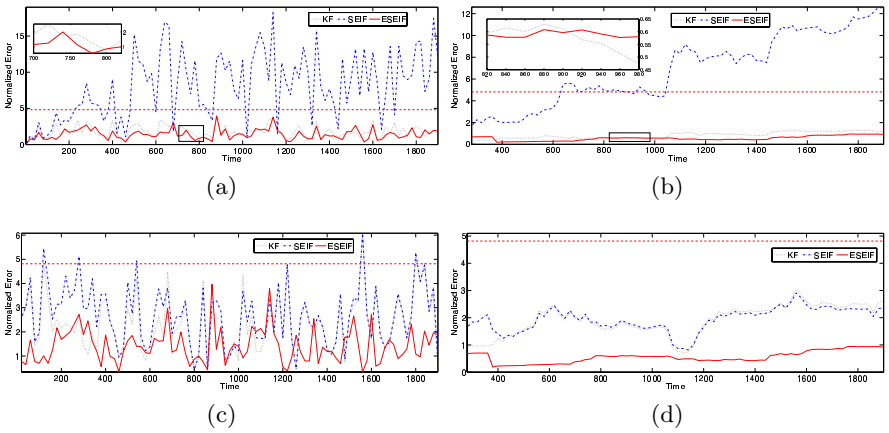

Fig. 3. The time history of the (a), (b) *global* and (c), (d) *local* normalized errors for the LG results, estimated from a series of Monte Carlo simulations. Plotted in (a) and (c) are the two errors for the vehicle. In (b) and (d) we show the errors for one of the features which is representative of the other elements in the map. The horizontal threshold denotes the $97.5^{\%}$ chi-square confidence bound. The local ESEIF and SEIF estimation errors are similar in magnitude to that of the Kalman Filter. The global error attributed to the SEIF, meanwhile, is noticeably larger, exceeding the chi-square bound. This indicates that the SEIF preserves local relationships but leads to estimates which are globally overconfident while the ESEIF maintains both global and local consistency.

the SEIF, on the other hand, are noticeably larger, frequently exceeding the chi-square bound. The local errors shown in Figures 3(c) and 3(d) are similar for all three filters, generally smaller than the confidence threshold. This behavior indicates that, in the linear Gaussian case, ESEIFs maintain a state estimate which is both globally and locally consistent while the SEIF leads to errors which are consistent locally but inconsistent in the absolute sense.

As a related consequence of the ESEIF sparsification strategy, the filter maintains conservative uncertainty estimates. In Figure 4(a) we compare the global map uncertainties for the two information filters to those of the Kalman Filter. In particular, from the inverse of the information matrices, we compute, for each feature, the log of the ratio of the covariance sub-block determinant to the determinant of the sub-block for the KF. Since the KF solution represents the true distribution, values larger than zero correspond to conservative estimates for a feature's position while values less than zero are a sign of overconfidence. As the histogram demonstrates, the ESEIF is conservative in its estimate for the absolute position of each feature while each of the marginals represented by the SEIF are overconfident. When we transform the maps relative to the first observed feature, we see in Figure 4(b) that the overconfidence of the SEIF is less severe while the ESEIF remains conservative. As a conse-

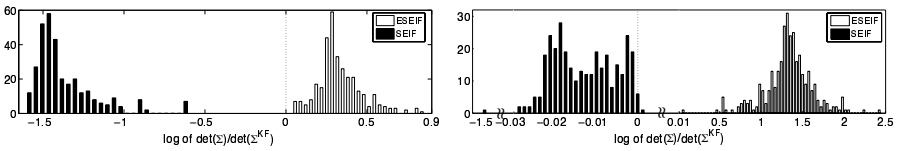

(a) (b)

Fig. 4. The LG simulation estimates of map uncertainty maintained by the ESEIF and SEIF compared with that of the KF. For each feature, we consider the log of the ratio of the covariance sub-block determinant for the information filters to the determinant for the KF. Values equal to zero indicate an exact estimate for the uncertainty. Log ratios greater than zero imply conservative estimates while values less than zero correspond to overconfidence. In (a) we show a histogram describing the global measure of uncertainty determined directly from the inverse of the information matrices. The SEIF yields map estimates which are largely overconfident while the ESEIF leads to estimates which are conservative. Depicted in (b), the overconfidence of the SEIF is less severe when we consider the relative map uncertainty which follows from root-shifting the state to the first feature added to the map. The one outlier corresponds to the original world origin as represented in the new reference frame. Meanwhile, the histogram shows that the ESEIF maintains conservative estimates for the relative map covariance matrix.

quence of the overconfidence of its global map, the one exception in the case of the SEIF is the representation of the original world origin in the root-shifted reference frame.

4.2 Experimental Validation

The linear Gaussian simulations allow us systematically analyze the accuracy of the sparsified filters when we are able to perform inference on an exact model. Unfortunately, for most real-world applications, both the vehicle motion and observation models are nonlinear and are corrupted by noise which is not Gaussian. To demonstrate the application of ESEIFs to typical SLAM problems, we implement the algorithm along with the SEIF and the EKF on two nonlinear datasets.

For the first real-world application of SLAM, we consider the benchmark Victoria Park dataset, widely used as a testbed for SLAM algorithms. A truck equipped with dead-reckoning sensors and a laser scanner drives in a series of loops within Victoria Park, Sydney. Using a simple perceptual grouping implementation, we are able to detect tree trunks located throughout the park among the laser data which is cluttered with spurious returns. We solve the data association problem offline to ensure that the correspondences are the same for each filter.

We implement the ESEIF and SEIF estimators together with the EKF which has been successfully applied to this dataset in the past. We limit

the number of active features to a maximum number of $\Gamma_a = 10$ for the information filters. When we perform sparsification in the ESEIF, our priority is again on relocation in that we reserve as many tree observations as possible (i.e. no more than $\Gamma_a = 10$) for the purpose of adding the vehicle back into the map. Remaining measurements, if any, are used to update the ESEIF prior to marginalization. This helps to minimize the influence of spurious data on the relocated vehicle pose.

We plot the ESEIF and SEIF estimates of the map together with the three sigma uncertainty bounds in Figures 5(a) and 5(b), respectively. The estimates of the 3 km trajectory for the car are superimposed on the plot. As a basis for comparison, the plots include the feature locations resulting from the EKF which are nearly identical to those published elsewhere. Both sparsified filters yield similar maps though the deviation from the EKF estimates is noticeably larger for the SEIF than it is for the ESEIF. Furthermore, the global confidence bounds for the ESEIF are conservative, yet comparable to the feature uncertainties maintained by the EKF while they are significantly overconfident for the SEIF. While not ground truth, the EKF represents the baseline which the information filters strive to match and, yet, many of the EKF estimates lie outside the three sigma uncertainty bounds for the SEIF. This is especially evident in the periphery as we indicate in the inset plot. As we saw in the LG simulation, all three algorithms seem to equivalently represent the local map relationships given by the transformation of the map into the vehicle's reference frame at its final pose. Both the ESEIF relative map shown in Figure 5(c) and the SEIF relative map in Figure 5(d) are almost identical to the corresponding EKF results. In this case, the relative ESEIF and SEIF uncertainty bounds now capture the EKF estimate for the feature locations. The SEIF algorithm allows us to achieve results which are similar to the standard EKF in the local but not global sense while ESEIFs provide a conservative map estimate which is nearly identical to the EKF both globally *and* locally.

We have seen from the plots of the two SLAM maps that SEIFs are much more confident in their state estimates. In Figure 6(a) we compare the global uncertainty of each feature for the ESEIF and SEIF to the EKF, again using the log of the ratio of the determinant of the feature covariances. As with the linear Gaussian simulations, the ESEIF log ratios are all greater than zero, indicating that ESEIFs maintain conservative estimates for the global uncertainty of each state element. On the other hand, those of the SEIF are largely overconfident. Expressing the state in the vehicle reference frame, the histogram in Figure 6(b) reveals that the SEIF remains overconfident, although to a lesser extent. The one exception is again the representation of the global origin in the vehicle frame and is a direct consequence of the global inconsistency of SEIFs. The ESEIF, meanwhile, remains conservative in the relative frame.

In the second experiment, a wheeled robot drives around a gymnasium in which 64 track hurdles are positioned at known locations along the baselines of

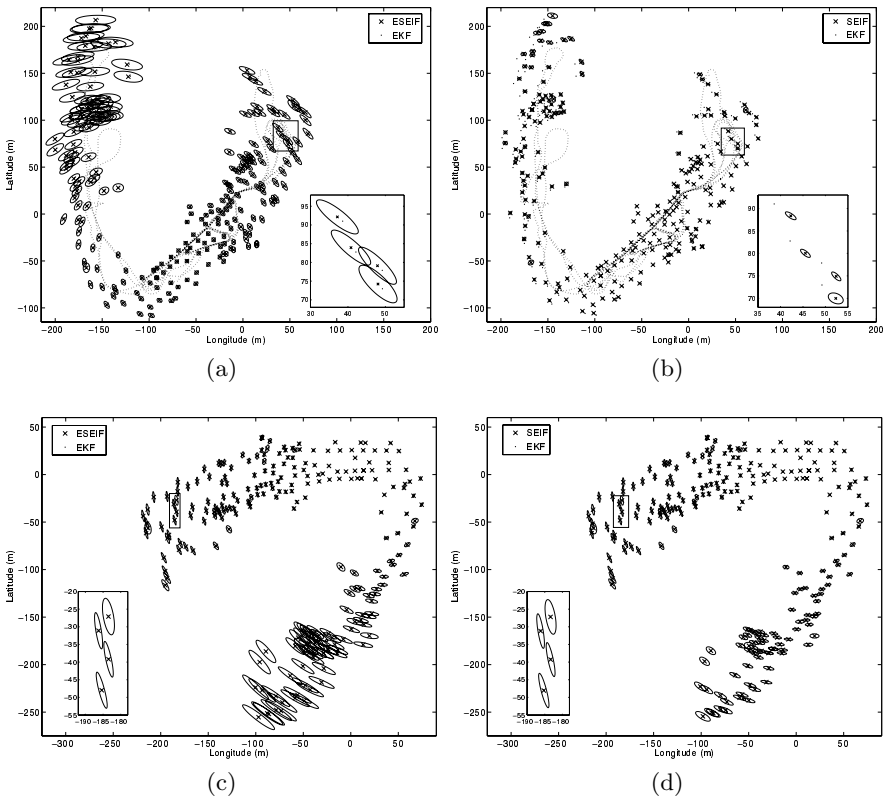

Fig. 5. Map and vehicle trajectory estimates for the Victoria Park dataset. In each, we include the final EKF map which agrees with previous results published in the literature. The top two plots represent the global state estimate while the two at the bottom are the result of root-shifting the map into the vehicle frame via compounding: $\mathbf{x}_{vi} = \ominus\mathbf{x}_v \oplus \mathbf{x}_i$. The plot in (a) presents the results of the ESEIF, including the three sigma confidence bounds for each of the features. The ESEIF produces feature estimates which are nearly identical to those of the EKF and, while it is omitted to make the plot readable, the uncertainty ellipses are very similar for the two filters. In (b), we see that while the SEIF and EKF maps are alike, the difference between the two estimates is noticeably larger for the SEIF algorithm. Additionally, the inset reveals that the SEIF yields global error estimates which are significantly overconfident. Looking at the maps expressed in the vehicle frame, though, we see that both (c) the ESEIF and (d) SEIF preserve the relative map structure.

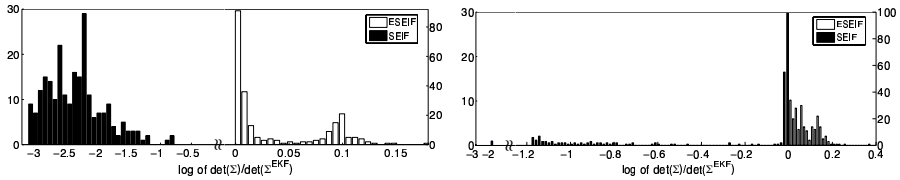

(a) (b)

Fig. 6. Histograms of the ESEIF and SEIF uncertainty estimates as compared to the EKF results for the Victoria Park dataset. We again use the log of the ratio of the covariance sub-block determinants for each feature. The histogram in (a) corresponds to the direct filter estimates and is representative of the global uncertainty. The ESEIF maintains conservative estimates for the uncertainties while the SEIF estimates are overconfident when compared to the EKF. Expressing the map in the vehicle's reference frame, (b) demonstrates that SEIFs remain overconfident but are better able to capture the relative uncertainty. Due to the global overconfidence, there is an outlier corresponding to the representation of the global origin in the robot's frame. Meanwhile, the ESEIF local estimates remain conservative relative to the EKF.

four adjacent tennis courts. Wheel encoders provide the input to the kinematic motion model while observations of the environment are made using a SICK laser scanner. Data association is again performed offline and is the same for each filter.

We perform SLAM on the data again using both the ESEIF and SEIF alongside a standard EKF implementation. When necessary, we employ the two sparsification strategies to maintain a bound of $\Gamma_a = 10$ active features. During ESEIF sparsification, we relocate the robot using a single feature observation which provides a measurement of the relative transformation (translation and rotation) between the vehicle and the hurdle.

In Figure 7(a), we show the final map estimated by the ESEIF, overlayed onto a depiction of the ground truth. The ellipses drawn around each feature correspond to the three sigma bound on the position of one of the hurdle legs. The same plot is shown in Figure 7(b) for the map estimated using the SEIF algorithm. Notice that the uncertainty bounds maintained by the SEIF are significantly overconfident and, for many hurdles, do not include the true feature position. While we are able to maintain an estimate of the state which is both globally *and* locally conservative compared with that of the EKF using ESEIFs, enforcing sparsity in the SEIF results in an estimate which suffers from global inconsistency.

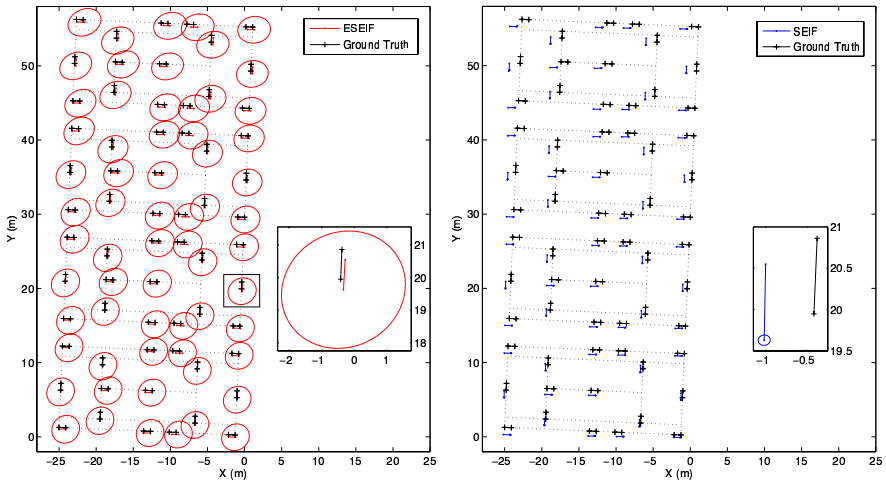

Fig. 7. The final maps generated by the (a) ESEIF and (b) SEIF algorithms. Included is an outline of the tennis courts overlayed with the ground truth hurdle poses indicated by the black cross hairs. The ellipses centered at the base leg of each feature represent the three sigma uncertainty bounds for position. Note the significant difference in magnitude between the confidence estimates maintained by the two filters. While the true feature locations are captured by the ESEIF uncertainty regions, a majority of the hurdles fall outside the SEIF ellipses. This overconfidence is a result of the approximation employed by SEIFs to enforce sparseness and is indicative of global inconsistency.

5 Discussion

We have shown both in simulation as well as with a pair of nonlinear datasets that the ESEIF maintains error measures which are both globally and locally conservative relative to the full Kalman estimates. In the linear Gaussian case, the implication is that the ESEIF sparsification strategy preserves consistency according to both metrics. On the other hand, as the ESEIF is formulated upon the dual of the EKF, it is subject to the same convergence issues that are attributed to the EKF for nonlinear applications [1]. As such, though the ESEIF error estimates are relatively conservative, this does not guarantee consistency in such cases. Nonetheless, the ESEIF algorithm is able to capitalize upon the computational benefits of the sparse information form without the cost of additional overconfidence. In this manner it provides an efficient means of achieving estimates nearly identical to those of the EKF which has been successfully applied in a number of real-world situations.

6 Conclusion

Of late, many researchers in the robotics community have been interested in developing solutions to the SLAM problem which scale with environments of arbitrary size. One approach that is particularly promising follows from the key insight that the information matrix is relatively sparse for feature-based SLAM. In the case where the matrix is exactly sparse, state estimation can be performed in near-constant time, irrespective of the number of landmarks in the environment.

While a majority of the elements in the information matrix are relatively weak, the matrix is naturally dense due to the effect of marginalizing out old robot poses. To achieve the efficiency benefits, the SEIF algorithm enforces sparsity by deliberately breaking weak links between the robot and the map. As a consequence of this pruning strategy, the SEIF state estimate suffers from global inconsistency.

In this paper, we have introduced an algorithm for feature-based SLAM which achieves an exactly sparse information matrix while maintaining global *and* local consistency, relative to the standard EKF. We have shown that, by periodically marginalizing out the robot and then relocalizing it within the map, we control the number of active landmarks and, in turn, the population of the information matrix. The ESEIF is then able to benefit from the efficiency of the sparse information form while yielding conservative estimates for the robot pose and map.

We have demonstrated the performance of ESEIFs, both in a systematic linear Gaussian simulation as well as on two different nonlinear datasets. In all three, we have shown that ESEIFs maintain estimates nearly identical to those of the EKF which, in comparison, are both globally *and* locally conservative.

Acknowledgements

This work was funded in part by the CenSSIS ERC of the NSF under grant EEC-9986821, by ONR under grants N00014-02-C-0210, N00014-05-1-0244, N00014-97-1-0202, and N00014-03-1-0879, and by the MIT Sea Grant College Program under grant NA86RG0074 (project RCM-3).

References

1. Y. Bar-Shalom, X. Rong Li, and T. Kirubarajan. *Estimation with Applications to Tracking and Navigation.* John Wiley & Sons, Inc., New York, 2001.
2. T. Duckett, S. Marshland, and J. Shapiro. Learning globally consistent maps by relaxation. In *Proceeding of the IEEE International Conference on Robotics and Automation (ICRA)*, pages 3841–3846, San Francisco, USA, 2000.

3. R. Eustice, H. Singh, and J. Leonard. Exactly sparse delayed state filters. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, pages 2428–2435, Barcelona, Spain, April 2005.
4. R. Eustice, H. Singh, J. Leonard, M. Walter, and R. Ballard. Visually navigating the RMS Titanic with SLAM information filters. In *Proceedings of Robotics: Science and Systems (RSS)*, Cambridge, MA, USA, June 2005.
5. R. Eustice, M. Walter, and J. Leonard. Sparse extended information filters: Insights into sparsification. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, Edmonton, Alberta, Canada, August 2005.
6. U. Frese. A proof for the approximate sparsity of SLAM information matrices. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, pages 331–337, Barcelona, Spain, April 2005.
7. U. Frese. Treemap: An $\mathcal{O}(\log n)$ algorithm for simultaneous localization and mapping. In C. Freksa, editor, *Spatial Cognition IV*, pages 455–476. Springer Verlag, 2005.
8. U. Frese and G. Hirzinger. Simultaneous localization and mapping - a discussion. In *Proceedings of the IJCAI Workshop on Reasoning with Uncertainty in Robotics*, pages 17–26, 2001.
9. U. Frese, P. Larsson, and T. Duckett. A multilevel relaxation algorithm for simultaneous localization and mapping. *IEEE Transactions on Robotics*, 21(2):196–207, April 2005.
10. Y. Liu and S. Thrun. Results for outdoor-SLAM using sparse extended information filters. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, pages 1227–1233, Taipei, Taiwan, 2003.
11. Peter S. Maybeck. *Stochastic Models, Estimation, and Control, Volume 1*. Academic Press, New York, NY, 1979.
12. M. Paskin. Thin junction tree filters for simultaneous localization and mapping. Technical Report CSD-02-1198, University of California, Berkeley, September 2002.
13. Jonathan Shewchuck. An introduction to the conjugate gradient method without the agonizing pain. Technical Report CMU-CS-94-125, Carnegie Mellon University, August 1994.
14. R. Smith, M. Self, and P. Cheeseman. Estimating uncertain spatial relationships in robotics. In I. Cox and G. Wilfong, editors, *Autonomous Robot Vehicles*, pages 167–193. Springer Verlag, 1990.
15. S. Thrun, Y. Liu, D. Koller, A.Y. Ng, Z. Ghahramani, and H. Durrant-Whyte. Simultaneous localization and mapping with sparse extended information filters. *International Journal of Robotics Research*, 23(7-8):693–716, July-August 2004.