

5-2008

Iterated Filters for Bearing-Only SLAM

Stephen Tully
Carnegie Mellon University

Hyungpil Moon
Sung Kyun Kwan University

George Kantor
Carnegie Mellon University

Howie Choset
Carnegie Mellon University

Follow this and additional works at: <http://repository.cmu.edu/robotics>

 Part of the [Robotics Commons](#)

Published In

Robotics and Automation, 2008. ICRA 2008. IEEE International Conference on , 1442-1448.

This Conference Proceeding is brought to you for free and open access by the School of Computer Science at Research Showcase @ CMU. It has been accepted for inclusion in Robotics Institute by an authorized administrator of Research Showcase @ CMU. For more information, please contact research-showcase@andrew.cmu.edu.

Iterated Filters for Bearing-Only SLAM

Stephen Tully, Hyungpil Moon, George Kantor, and Howie Choset

Abstract—This paper discusses the importance of iteration when performing the measurement update step for the problem of bearing-only SLAM. We focus on an undelayed approach that initializes a landmark after only one bearing measurement. Traditionally, the extended Kalman filter (EKF) has been used for SLAM, but the EKF measurement update rule can often lead to a divergent state estimate due to its inconsistency in linearization. We discuss the flaws of the EKF in this paper, and show that even the well established inverse-depth parametrization for bearing-only SLAM can be affected. We then show that representing the bearing-only update as a numerical optimization problem (solved with an iterative approach such as Gauss-Newton minimization) prevents divergence of the Kalman filter state and produces accurate SLAM results for a bearing-only sensor. More specifically, we propose the use of an iterated Kalman filter to resolve the issues normally associated with the EKF measurement update. Two outdoor mobile robot experiments are discussed to compare algorithm performance.

I. INTRODUCTION

Simultaneous localization and mapping (SLAM) is the task of incrementally building a map of the surrounding environment with a mobile robot while simultaneously localizing the robot in the map. With the availability of inexpensive monocular vision, bearing-only SLAM has received increased attention recently [1]–[8]. Many of these techniques accumulate bearing measurements until a new landmark can be accurately placed in the environment. Instead, we focus on an undelayed approach that arbitrarily initializes observed landmarks. With undelayed initialization, the Kalman filter can immediately gain the localization benefit of having measured a new landmark.

A common tool used for SLAM is the extended Kalman filter (EKF). Unfortunately, the extended Kalman filter is ill-suited for the problem of bearing-only SLAM because of a highly nonlinear measurement model. In many situations, the state estimate will diverge, resulting in an inaccurate SLAM result.

We address this problem by observing that the EKF measurement update is equivalent to performing a single iteration of Gauss-Newton optimization on the cost function that is defined for the Kalman update [9]. When performing numerical optimization, it does not make sense to take just one step in the Newton direction. The update equation should be iterated until convergence.

S. Tully is with the Electrical and Computer Engineering Department and G. Kantor, H. Choset are with the Robotics Institute at Carnegie Mellon University, Pittsburgh, PA 15213, USA. {stully@ece, kantor@ri, choset@cs}.cmu.edu

H. Moon is with the Mechanical Engineering Department at SungKyunKwan University, Suwon, Gyeonggi-do 440-746, Korea. hyungpil@me.skku.ac.kr




Fig. 1. An illustration of the measurement update task for bearing-only SLAM.

We adopt the iterated Kalman filter (IKF) for bearing-only SLAM. It can be shown that the IKF is equivalent to using a Gauss-Newton method for optimizing the Kalman update [9]. In some literature (for example in [10]), it is claimed that the IKF update may enhance divergence. However, we show that when implemented properly with a variable-step *backtracking* method as described in [11], the IKF will in fact counteract divergence. Although the iterated form of the Kalman filter is not new, we believe we are the first to properly demonstrate its use with bearing-only SLAM.

Recent work suggests encoding landmarks with a modified polar coordinate representation that estimates inverse-depth to landmarks [8], [12]. This representation results in a measurement equation that has low linearization error. We show that, even for this improved parameterization, the extended Kalman filter can fail. Iterating the measurement update, as we are trying to emphasize throughout this paper, improves performance.

In this paper, we briefly review other methods in Sec. II and then formalize in detail the filtering task of bearing-only SLAM, including the definitions of two different landmark parameterizations, in Sec. III. We then demonstrate the ill effects of relying on the EKF in Sec. IV in order to motivate our need for an iterated filter, which we develop fully in Sec. V. Finally, we introduce experimental evidence that demonstrates the benefit of iteration for bearing-only SLAM in Sec. VI before concluding the paper with a discussion in Sec. VII.

II. RELATED WORK

Delaying the initialization of landmarks is a popular method in bearing-only SLAM. A batch update with all of the stored observations is demonstrated in [13]. In [1], initialization is postponed until a pair of measurements are distinguishable enough and the probability density of the corresponding landmark becomes sufficiently Gaussian. In [2], the persistence of landmark pose estimation is tracked without prior knowledge of data association. To incorporate the positioning and sensing uncertainties, the authors project measurements from the sensor space to the plane by approximating Gaussian distributions with bivariate ellipse representations.

Another popular method, which is related to the extended Kalman filter, is the Gaussian sum filter (GSF). The GSF approximates arbitrary probability density functions by a weighted combination of many multivariate Gaussians. Since the GSF, which is introduced in [14], [15], requires retaining a large set of EKFs, its computational complexity can become problematic when a large number of landmarks are initialized simultaneously (typical for vision-based methods). An approximated Gaussian sum method is proposed in [4], where a set of parametrized cascaded Gaussian distributions and a single covariance matrix for all the Gaussians is managed and updated by federated information sharing (FIS). Although this method is an improvement over the standard GSF, a larger state is still required to initialize landmarks.

Particle Filters (PFs), which incorporate non-Gaussian distributions, are widely used in SLAM research. In [5], particle filters are adopted for bearing measurements by associating hypothesized pseudo-ranges with each bearing measurement and by implementing a re-sampling procedure to eliminate improbable particles. In [16], a set of particles are maintained along the viewing ray of a landmark and landmark initialization is delayed until the range distribution is roughly Gaussian. In [7], a FastSLAM particle filter is used for single-camera SLAM with a partial initialization strategy which estimates the inverse-depth of new landmarks rather than their depth. Unfortunately, particle filter methods often require a very large number of particles for landmark initialization.

Direct parametrization of inverse-depth is used for monocular SLAM in [12] and a method of inverse-depth and depth conversion is proposed by the same authors in [8]. It is well known that the inverse-depth representation can handle distant landmarks more efficiently than conventional parameterizations because it incorporates a measurement equation with low linearization error. This method effectively doubles the size of the Kalman filter state in order to encode landmark locations, which increases the computational complexity of the problem.

III. THE BEARING-ONLY SLAM PROBLEM

In this section, we will present two formalizations of the undelayed bearing-only Kalman filter SLAM problem. The first is more conventional and encodes the location of a

landmark by its x and y coordinates. The second uses an inverse-depth parameterization, as presented in [12].

A. Defining the Kalman Filter State

The conventional (x,y) definition of the filter state is

$$X_k = [x_R \ y_R \ \theta_R \ x_{L_1} \ y_{L_1} \ x_{L_2} \ y_{L_2} \ \dots \ x_{L_N} \ y_{L_N}]^T,$$

which represents the robot pose appended by the locations of N observed landmarks. The inverse-depth parameterization proposed in [12] modifies the state as follows.

$$X_k = [x_R \ y_R \ \theta_R \ x_{L_1}^0 \ y_{L_1}^0 \ \rho_{L_1} \ \phi_{L_1} \ \dots \ x_{L_N}^0 \ y_{L_N}^0 \ \rho_{L_N} \ \phi_{L_N}]^T,$$

where $x_{L_i}^0$ and $y_{L_i}^0$ are the x and y coordinates of the robot when the i th landmark is initialized, and ρ_{L_i} and ϕ_{L_i} are the inverse-depth and bearing to the i th landmark relative to that initialization point. The distribution over possible states is estimated recursively by its mean \hat{X}_k and covariance P_k .

B. Models for Bearing-Only SLAM

The motion input, $u_k = [v_k \ \omega_k]^T$, contains the translational and rotational velocities of the robot at time-step k . The state evolves according to the process model

$$f(X_k, u_k) = X_k + \begin{bmatrix} v_k \cos(\theta_R) \Delta t \\ v_k \sin(\theta_R) \Delta t \\ \omega_k \Delta t \\ 0_{2N,1} \end{bmatrix}$$

where Δt is the time elapsed since the previous update and U is the covariance of the input u_k .

The (x,y) measurement model for a relative bearing observation of the i th landmark is

$$h_i(X_k) = \arctan\left(\frac{y_{L_i} - y_R}{x_{L_i} - x_R}\right) - \theta_R.$$

The inverse-depth equivalent is

$$h_i(X_k) = \arctan\left(\frac{y_{L_i}^0 + \frac{1}{\rho_{L_i}} \sin \phi_{L_i} - y_R}{x_{L_i}^0 + \frac{1}{\rho_{L_i}} \cos \phi_{L_i} - x_R}\right) - \theta_R.$$

We assume that all bearing measurements are perturbed by additive white Gaussian noise with variance σ_z^2 .

C. Prediction Step

For the Kalman filter prediction step, the state mean and covariance matrix are computed as follows.

$$\begin{aligned} \hat{X}_{k+1}^- &= f(\hat{X}_k^+, u_k) \\ P_{k+1}^- &= F P_k^+ F^T + W U W^T \end{aligned}$$

where \hat{X}_k^+ and P_k^+ define the estimate from the previous time step, and F and W are the Jacobians of the state process equation.

$$F = \frac{\partial f(\hat{X}_k^+, u_k)}{\partial X_k}, \quad W = \frac{\partial f(\hat{X}_k^+, u_k)}{\partial u_k}.$$

D. Landmark Initialization

In order to satisfy true undelayed estimation, a landmark must be initialized after observing only one relative bearing measurement. We will refer to this first observation as z_{init} . Without any information regarding the range of the landmark along this initial measurement ray, we must guess or use some heuristic for choosing r_{init} , the initialization range.

For the conventional (x,y) filter, the initialized landmark location can be computed by mapping the relative bearing measurement z_{init} and the chosen initialization range r_{init} through a nonlinear function

$$g(X_k, z_{init}, r_{init}) = \begin{bmatrix} x_R + r_{init} \cos(\theta_R + z_{init}) \\ y_R + r_{init} \sin(\theta_R + z_{init}) \end{bmatrix}.$$

The analogous function for encoding a new landmark with the inverse-depth parameterization is

$$g(X_k, z_{init}, r_{init}) = \begin{bmatrix} x_R \\ y_R \\ \frac{1}{r_{init}} \\ \theta_R + z_{init} \end{bmatrix}.$$

The state mean and covariance are augmented as follows to include the new landmark in the Kalman filter state.

$$\hat{X}_k^- = \begin{bmatrix} \hat{X}_k^- \\ g(\hat{X}_k^-, z_{init}, r_{init}) \end{bmatrix}$$

$$P_k^- = \begin{bmatrix} P_k^- & P_k^- \frac{\partial g}{\partial X_k^- T} \\ \frac{\partial g}{\partial X_k^-} P_k^- & \frac{\partial g}{\partial X_k^-} P_k^- \frac{\partial g}{\partial X_k^- T} + \frac{\partial g}{\partial v} V \frac{\partial g}{\partial v T} \end{bmatrix}$$

where $v = [r_{init} \ z_{init}]^T$ for the conventional (x,y) filter and $v = [\frac{1}{r_{init}} \ z_{init}]^T$ for the inverse-depth parameterization. V is the covariance matrix for the initialization vector v , which is chosen by the user. For our experiments, we define the initialization covariance

$$V_{xy} = \begin{bmatrix} A & 0 \\ 0 & \sigma_z^2 \end{bmatrix}, \quad V_{ID} = \begin{bmatrix} 1.0 & 0 \\ 0 & \sigma_z^2 \end{bmatrix}$$

for the (x,y) filter and the inverse depth filter respectively, with A arbitrarily large.

E. Measurement Update

The measurement update is the central topic of this paper. The purpose of the measurement update is to improve our state estimate by incorporating bearing observations of landmarks already present in the Kalman filter state. To motivate the use of an iterated method, we will review the overall objective of the Kalman update here (a task that is illustrated in Fig. 1).

Ideally, we would like to replace the current state mean with the maximum likelihood state estimate given the measurement z_k and the predicted state estimate \hat{X}_k^-, P_k^- . This is equivalent to maximizing the following posterior probability.

$$\hat{X}_k^+ = \arg \max_X \text{prob}(X | z_k, \hat{X}_k^-, P_k^-) \quad (1)$$

where z_k is a vector composed of all bearing observations acquired at time-step k .




Fig. 2. An analytical example where a robot with perfect odometry and a perfect bearing sensor acquires two measurements to a landmark at the origin (filled star). The EKF update fails to localize the landmark at the origin when the initialization is incorrect (open star), despite perfect measurements.

When adopting the framework of the Kalman filter, the solution to (1) is the minimization of the following nonlinear least-squares cost function

$$c(X) = \begin{bmatrix} z_k - h(X) \\ X - \hat{X}_k^- \end{bmatrix}^T \begin{bmatrix} R & 0 \\ 0 & P_k^- \end{bmatrix}^{-1} \begin{bmatrix} z_k - h(X) \\ X - \hat{X}_k^- \end{bmatrix} \quad (2)$$

where R is a matrix with diagonal elements equal to σ_z^2 , the variance of a single bearing measurement.

Given a linear measurement model, the standard Kalman filter measurement update equation would provide a closed-form optimal solution to the minimization of (2). This is not the case with bearing-only SLAM.

The extended Kalman filter measurement update approximates the solution of (2) via linearization of the measurement function about the current estimate. In the next section we will show that this option will not work well with bearing-only measurements.

IV. FAILURES OF THE EKF

The widely used extended Kalman filter measurement update can be written in the following way.

$$H = \frac{\partial h(\hat{X}_k^-)}{\partial X_k}$$

$$\hat{X}_k^+ = \hat{X}_k^- + P_k^- H^T (H P_k^- H^T + R)^{-1} (z_k - h(\hat{X}_k^-))$$

$$P_k^+ = P_k^- - P_k^- H^T (H P_k^- H^T + R)^{-1} H P_k^- \quad (3)$$

where X_k^- is the previous state mean, P_k^- is the previous covariance matrix, and H is the Jacobian of the measurement function linearized about the current estimate. It is due to this linearization that problems arise when using the EKF measurement update.

We will now work through an analytical exercise that demonstrates the undesirable effects of the EKF. Let us assume we have a mobile robot, with perfect odometry, driving along a path defined by a circle of radius one centered at the origin, as shown in Fig. 2. Also, the robot can measure a landmark at the origin with a perfect bearing sensor. In

Fig. 2, we show the robot observing the landmark once at a position of $(-1,0)$ and again at a position of $(0,1)$. Two such measurements, from two distinctive and known positions, should be enough to localize the exact location of the landmark regardless of the parameterization used to encode landmarks. We believe that a decent nonlinear estimator should, at the very least, solve this simplified example by finding the maximum likelihood estimate.




Fig. 3. The resulting landmark location versus its initialized location for an example in which two perfect bearing measurements are obtained from distinctive locations. Top: the (x,y) filter result. Bottom: the inverse-depth result.

When deriving the (x,y) EKF update result for the analytical example illustrated in Fig. 2, we end up with

$$x_1 = x_0 - (x_0^2 + 1) \arctan(x_0) \quad (4)$$

where x_0 is the initialized location of the landmark on the x -axis (equal to $r_{init} - 1$), and x_1 is the resulting estimate of the landmark location on the x -axis after two perfect bearing measurements. Notice that if $x_0 \neq 0$, x_1 will not equal the correct value. We note that the y coordinate of the landmark is correctly localized for this example. Fig. 3 shows the resulting x -axis location of the landmark estimate versus different initialization points.

The previous result shows that, unless you can somehow exactly guess the depth of the landmark after just one bearing measurement, the estimate will be wrong. It is also worth noting that, when the landmark is initialized at a large depth, the EKF update will tend to over-correct and place the estimate to the left of the original robot position. This is an undesirable result because the new location will no longer agree with the first measurement. Often this problem may lead to divergence after additional measurements.

Although it is true that using inverse-depth parametrization makes the bearing only SLAM problem more linear for low parallax situations, it is not necessarily an improvement for high parallax measurements. The nonlinearity issues pertaining to the EKF measurement update still remain. After

carrying out the EKF update with inverse-depth parameterization and then mapping the result back to an (x,y) coordinate, the following update rule is obtained for the previous analytical exercise.

$$x_1 = \frac{(x_0 + 1)^2}{x_0 + 1 + (x_0^2 + 1) \arctan(x_0)} - 1.$$

In this case, the updated landmark estimate for various initialization positions is shown in Fig. 3.

It is clear that, under certain unlucky choices of x_0 , the updated landmark estimate would diverge to a negative value very far away from the robot and to the left of its original position. This demonstrates the possibility of a very large error despite two perfect measurements. This situation involves replacing the current inverse-depth estimate with a negative value. Under the assumptions of this parameterization technique, we should not have a negative depth, and in fact such an occurrence causes divergence in experiments.

V. AN ITERATED FILTER FOR BEARING-ONLY SLAM

As discussed in Sec. III, the objective of the Kalman measurement update is to find the maximum likelihood state estimate by minimizing the cost function given by (2). We have previously discussed that the extended Kalman filter approximates this minimization through linearization about the current estimate, but in many cases does not solve the true minimum of (2). We propose the use of numerical optimization to minimize the objective cost function because of the nonlinearity of bearing-only measurement models.

The Gauss-Newton algorithm iteratively solves the nonlinear least-squares minimization problem with the following recursive equation.

$$X_{i+1} = X_i - \gamma_i (\nabla^2 c(X_i))^{-1} \nabla c(X_i) \quad (5)$$

where γ_i is a parameter to vary the length of the step-size.

By substituting the appropriate Jacobian $\nabla c(X_i)$ into (5) along with the Hessian $\nabla^2 c(X_i)$, we arrive at the following update rule for the iterated Kalman filter, where H is the Jacobian of the measurement function $h(X)$ linearized about X_i . We refer the reader to [9] for details of this derivation.

$$\begin{aligned} x_0 &= \hat{x}_{k+1}^-, P_0 = P_{k+1}^- \\ K_i &= P_0 H_i^T (H_i P_0 H_i^T + R)^{-1} \\ x_{i+1} &= x_i + \gamma_i (H_i^T R^{-1} H_i + P_0^{-1})^{-1} \\ &\quad \cdot (H_i^T R^{-1} (z - h(x_i)) + P_0^{-1} (x_0 - x_i)) \quad (6) \\ &= x_0 + K_i (z_{k+1} - h(x_i) - H_i (x_0 - x_i)) \quad (7) \end{aligned}$$

In the equations above, the subscript k refers to the time-index and the subscript i refers to the iteration index. For a single update step at time k there may require many iterations of i before the state x_i converges. Once the IKF does converge, the estimate is overwritten with the result, $\hat{X}_{k+1}^+ = X_i$. In order to obtain the conventional IKF update equation, as in 7, the variable-step parameter γ_i in (6) must equal one.

We believe it is important to recognize that the EKF update rule is equivalent to performing just one iteration of




Fig. 4. A vision-based experimental result of bearing-only SLAM overlapped with a satellite image. This map was created using the (x,y) IKF method. A cyan (light colored) solid line depicts the robot trajectory, uncertainty ellipses are shown for the map estimate, and \times marks depict the true landmark positions.

the Gauss-Newton algorithm: taking (7) and setting $i = 0$ produces the well known EKF update equations below.

$$\begin{aligned} K_0 &= P_0 H_0^T (H_0 P_0 H_0^T + R)^{-1} \\ \hat{x}_{k+1}^+ &= x_1 = x_0 + K_0 (z_{k+1} - h(x_0)) \end{aligned}$$

Typical formulations of the IKF update equation, for example in [17], always choose γ_i equal to one and therefore force a full Newton step when optimizing. For the problem of bearing-only SLAM, taking a full Newton step is risky because it may increase the cost and lead to divergence of the state estimate.

We incorporate *backtracking* and revert to (6). By varying γ_i to satisfy an acceptance criteria [11], we instead take a partial step in the Newton direction and can guarantee that the new estimate computed during iteration i reduces the cost.

In our experiments, the optimization requires between two and five iterations to converge, and is therefore only slightly more computationally intensive than the conventional EKF. On the other hand, we must reevaluate γ_i several times within one iteration step of (6) to ensure that the cost is reduced, but this is not processor intensive because the correction term of (6) can be precomputed and simply scaled. Although we iterate to improve our estimate for the state mean, we use the same covariance update method that is used with the conventional EKF measurement update (3).

VI. EXPERIMENTAL RESULTS

We have applied our iterated filter to both parameterization methods described in Sec. III and to two different data sets. One is an outdoor vision-based experiment and another is based on the standard Victoria Park benchmark dataset [18].




Fig. 5. This plot shows the accuracy of the completed map for a visual SLAM experiment while varying the value of r_{init} . This plot compares the conventional x,y EKF (labeled EKF), the x,y IKF (labeled IKF), the inverse-depth EKF (labeled ID), and the inverse-depth IKF (labeled IID).




Fig. 6. This plot shows the accuracy of the completed map for a visual SLAM experiment while varying the measurement frequency. This plot compares the conventional x,y EKF (labeled EKF), the x,y IKF (labeled IKF), the inverse-depth EKF (labeled ID), and the inverse-depth IKF (labeled IID).

For each of the experiments, we assume data association has been solved.

The first experiment is an offline vision-based SLAM experiment (see Fig. 4) with a robot that captures images at 15 Hz with four cameras that have non-overlapping fields of view. Although this experiment does not involve a large number of landmarks, it offers other challenges for bearing-only SLAM such as low parallax observations and high parallax observations. One of the landmarks in the map is a distance of 100m from the robot when it is initialized.

In Fig. 5, we show the accuracy of the experiment while varying the initialization range for four different estimation methods: the conventional (x,y) EKF, the (x,y) IKF, the inverse-depth EKF, and the inverse-depth IKF. The conventional (x,y) EKF result is only partially represented in this plot because choosing a large initialization range causes this method to diverge. On the other hand, the (x,y) IKF result prevents divergence for any initialization depth and produces




Fig. 7. The EKF inverse-depth method must ignore a measurement update that results in a negative range, effectively “blinding” the robot for a period of time. This shows divergence of the robot pose.

a very accurate map that outperforms the inverse-depth methods in many cases. Additionally, this figure exemplifies the benefit of the inverse-depth parameterization for landmark initialization: the accuracy of the map is completely independent of the chosen initialization depth. Because, at a 15 Hz measurement rate, all observations for our data is effectively low parallax, neither inverse-depth algorithm was effected by the nonlinearity issues discussed in Sec. IV.

Fig. 6 shows the accuracy of the map while varying the effective measurement update rate. The reason we are looking at this relationship is because decreasing the measurement rate introduces high parallax observations. High parallax situations can occur when landmarks are very close, when the robot is travelling very fast, when the measurement rate is slowed by image processing, and/or when occlusion causes measurements to originate from very different vantage points. Fig. 6 shows that both the conventional (x,y) EKF and the inverse-depth EKF bearing-only SLAM algorithms are effected when the measurement rate is decreased, therefore implying that these algorithms are susceptible to high parallax situations. Both iterated methods work well throughout this test and have comparable accuracy. A typical result for a map created with the (x,y) IKF is shown in Fig. 4 and shows proper convergence of landmark locations despite the nonlinearity of the measurement model.

The next experiment is extracted from the popular Victoria Park SLAM dataset [18], which offers logged range/bearing measurements from a laser range sensor. To conform to the bearing-only theme of this paper, we appropriately discard all range measurements offered within this dataset. Although the landmarks are nearby and therefore do not offer a challenging low parallax test for the presented algorithms, there are numerous landmarks. We have found that the failures of the EKF emerge more quickly for a larger dataset.

In Sec. IV we discuss the possibility that the inverse-depth EKF update can result in a negative inverse-depth value. This can be very problematic and, in fact, a strict implementation of the inverse-depth EKF will quickly result in a diverging state estimate during the first several seconds



Fig. 8. An experimental result of bearing only SLAM using the Victoria Park dataset. A solid line depicts the robot trajectory and yellow dots depict landmarks.

of the Victoria Park dataset due to this reason. An easy solution is to effectively ignore a measurement update if it results in a negative inverse-depth value (because the clearly erroneous update would cause more harm than simply reinstating the previous estimate). Unfortunately, as seen in Fig. 7, this implementation hack cannot save the inverse-depth EKF method. Throughout the path drawn in red (darker solid line), the robot is effectively blind because it continues to ignore a sequence of 46 consecutive measurement updates that attempt to encode a landmark with a negative inverse-depth. Finally the estimate is eventually corrected towards the proper estimate, drawn in green (lighter solid line). Ignoring measurement updates is not a good solution to this problem: in fact, over 33 percent of measurements are discarded throughout the portion of the Victoria Park experiment we used to compare our algorithms. An iterated method, on the other hand, is able to apply every measurement correctly and produces a very accurate map.

The conventional (x,y) EKF is understandably a disaster when applied to the Victoria Park dataset. The state estimate diverges quickly, moving the robot a significant distance away from the proper pose location and converging landmark estimates in completely wrong places. The (x,y) IKF, on the other hand, which we formalized in Sec. V, has no problems estimating a large number of landmarks in this large-scale experiment, as seen in Fig. 8.

VII. CONCLUSION

Throughout this paper, we have ignored the problem of data association. Although data association is in itself a major topic in the robotics community, we chose to focus more on the estimation aspects of this problem with the assumption that data association can be solved through other methods (such as feature matching based on visual information). We also demonstrated our algorithms through offline experiments. This was to aid in the task of parameter tuning, and we believe that all algorithms in this paper would

easily run online (with a limit on the number of landmarks in the state).

For the most part, SLAM researchers have opted to use multi-hypothesis filters, delayed initialization, or a different parameterization of the state to solve the problem of undelayed bearing-only SLAM. The reason is that the measurement model for bearing-only observations is highly nonlinear, which causes the conventional extended Kalman filter measurement update to diverge for large linearization errors.

One contribution we have made is to show that the traditional Kalman filter framework used in numerous other SLAM papers (where landmarks are encoded by their x,y location) can be adopted for undelayed bearing-only SLAM if properly implemented with an iterated measurement update. Although the iterated Kalman filter is not new, we believe we are the first to apply it to the problem of bearing-only SLAM. The (x,y) IKF method works well for outdoor experiments that include far away landmarks and low parallax situations.

One advantage of using an inverse-depth parameterization is its ability to handle extreme low parallax observations. But in our paper, we point out several situations in which the inverse-depth method can run into problems. In this case, an iterated form of the Kalman filter is again the solution.

Our main contribution is to show that iterated methods are beneficial for bearing-only SLAM. The parameterization to use is, in our opinion, a choice of the user. If a larger state size becomes problematic then the (x,y) IKF formalization may be preferred. If robustness to extreme low parallax situations is necessary then the iterated inverse depth filter would be more appropriate. Even though the proper choice may depend on the specific application, our main point still remains: iteration is crucial for accurate bearing-only SLAM.

In this paper, we apply Gauss-Newton method for numerical optimization. The purpose of iteration is to minimize the objective cost function. Other methods of numerical optimization would likely produce equivalent results, such as the Broyden method or the Levenberg-Marquardt algorithm. We are currently investigating these different optimization methods. Additionally, although we have been unable to prove that numerical optimization of the cost function introduced in Sec. III will always converge upon the *global* minimum, we believe it to be true (via extensive experimental evidence). If this is indeed true, then the iterated form of the Kalman filter (for this specific problem) would produce the optimal state mean update at every time-step. This, in no way, suggests that the algorithms presented here would be optimal filters. The covariance update is still based on linearization and is an approximation. In future work we will try to prove the optimality of the update and are currently investigating more appropriate covariance update methods to improve performance.

ACKNOWLEDGMENTS

We would like to thank Jose Guivant, Juan Nieto and Eduardo Nebot for providing the Victoria Park dataset.

REFERENCES

- [1] T. Bailey, "Constrained initialisation for bearing-only SLAM," in *Proceedings of the 2003 IEEE International Conference on Robotics and Automation*, September 2003.
- [2] A. Costa, G. Kantor, and H. Choset, "Bearing-only landmark initialization with unknown data association," in *Proceedings of the 2004 IEEE International Conference on Robotics and Automation*, April 2004.
- [3] N. M. Kwok, G. Dissanayake, and Q. P. Ha, "Bearing-only SLAM using a SPRT based gaussian sum filter," in *Proceedings of the 2005 IEEE International Conference on Robotics and Automation*, April 2005.
- [4] J. Solà, A. Monin, M. Devy, and T. Lemaire, "Undelayed initialization in bearing only SLAM," in *Proceedings of the 2005 IEEE International Conference on Robotics and Automation*, April 2005.
- [5] N. M. Kwok and A. B. Rad, "A modified particle filter for simultaneous localization and mapping," *Journal of Intelligent and Robotic Systems*, vol. 46, no. 4, pp. 365–382, 2006.
- [6] K. E. Bekris, M. Glick, and L. Kavraki, "Evaluation of algorithms for bearing-only SLAM," in *Proceedings of the 2006 IEEE International Conference on Robotics and Automation*, May 2006.
- [7] E. Eade and T. Drummond, "Scalable monocular slam," in *IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, June 2006.
- [8] J. Civera, A. J. Davison, and J. M. M. Montiel, "Inverse depth to depth conversion for monocular slam," in *Proceedings of the 2007 IEEE International Conference on Robotics and Automation*, April 2007.
- [9] B. Bell and F. Cathey, "The iterated kalman filter update as a gauss-newton method," *IEEE Transactions on Automatic Control*, vol. 38, no. 2, pp. 294–297, 1993.
- [10] T. Bailey, J. Nieto, J. Guivant, M. Stevens, and E. Nebot, "Consistency of the ekf-slam algorithm," in *Proceedings of IEEE/RSJ International Conference on Intelligent Robots and Systems*, October 2006.
- [11] W. H. Press, S. A. Teukolsky, W. T. Vetterling, and B. P. Flannery, *Numerical Recipes in C*, 2nd ed. Cambridge University Press, 1992.
- [12] J. M. M. Montiel, J. Civera, and A. J. Davison, "Unified inverse depth parametrization for monocular slam," in *Robotics Science and Systems*, 2006. [Online]. Available: <http://www.roboticsproceedings.org/rss02/p11.html>
- [13] M. Deans and M. Hebert, "Experimental comparison of techniques for localization and mapping using a bearings only sensor," in *Proceedings of the 7th International Symposium on Experimental Robotics*, December 2000.
- [14] H. W. Sorenson and D. L. Alspach, "Recursive bayesian estimation using gaussian sums," *Automatica*, vol. 7, pp. 465–479, 1971.
- [15] D. L. Alspach and H. W. Sorenson, "Nonlinear bayesian estimation using gaussian sum approximations," *IEEE Transactions on Automatic Control*, vol. 17, no. 4, pp. 439–448, 1972.
- [16] A. Davison, "Real time simultaneous localisation and mapping with a single camera," in *International Conference on Computer Vision*, July 2003.
- [17] Y. Bar-Shalom, X. R. Li, and T. Kirubarajan, *Estimation with applications to tracking and navigation*. John Wiley and Sons, 2001.
- [18] J. Guivant, J. Nieto, and E. Nebot, Victoria park dataset. [Online]. Available: www.acfr.usyd.edu.au/homepages/academic/enbot/dataset.htm