Proceedings of the 2009 IEEE International Conference on Systems, Man, and Cybernetics
San Antonio, TX, USA - October 2009

Simultaneous Localization and Mapping in Dense Environments

S. Chakravorty, R. Saha
Department of Aerospace Engineering
Texas A&M University, College Station

***Abstract*— A hybrid Bayesian/ frequentist approach is presented for the Simultaneous Localization and Mapping Problem (SLAM). A frequentist approach is proposed for mapping with time varying robotic poses and is generalized to the case when the robotic pose is both time varying and uncertain. The SLAM problem is then solved in two steps: 1) the robot is localized with respect to a sparse set of landmarks in the map using a Bayes filter and a belief on the robot pose is formed, and 2) this belief on the robot pose is used to map the rest of the map using the frequentist estimator. The hybrid methodology is shown to have complexity linear in the map components, is robust to the data association problem and is provably consistent.**

I. Introduction

The problem of simultaneous, localization and mapping is considered in this paper. In the proposed method, the environment is split into a small set of sparse landmarks/ features and the rest of the (dense) environment. The philosophy behind the hybrid method amounts to the following two steps: 1) localize with respect to the landmarks, i.e., form a belief on the pose of the robotic system based on observations of the landmarks using a Bayes filter such as an EKF(**the Bayesian sub-problem**), and then, 2) map the rest of the environment based on the belief on the robotic pose using the frequentist mapping technique (**the frequentist sub-problem**). This formulation has linear complexity in the map components, is robust to the data association problem and provably consistent.

There are two main categories of approaches to SLAM: recursive and trajectory based.
In the recursive Kalman filter/ Information filter based approach [1]–[5], the map is appended to the filter as a parameter and the joint pose-map pdf is estimated using the Kalman recursion in either the covariance or the information matrix form. The Kalman filter scales quadratically as the size of the environment since the correlations between all the map elements need to be maintained in order for the map to be consistent [1] owing to the nature of the SLAM problem. The EKF SLAM based approach can only extract sparse maps since it is not robust to the data association problem in dense maps. In the SEIF (sparse extended information filter) based approach, the information form of the Kalman update is used and the sparsity of the information filter used to obtain nearly contant time SLAM algorithms [4]. However, the SEIf based filter tends to get overconfident and the latest research has concentrated on alleviating this problem by either having exact sparsity by keeping track of the trajectory [5] or enforcing sparsity by solving an approximate problem involving "kidnapping and relocalizing" the robot [6]. In the approach outlined in this paper, due to the frequentist update of the map, the correlations between the various map elements need not be maintained for consistency, which is in contrast to the result in [1] which asserts that maintaining correlations is key to consistency. The difference is due to the frequentist problem formulation which obviates the need to keep track of the correlations between the various map components. In addition, it is also shown that the hybrid formulation is robust to the data association problem in dense maps. In our opinion, the hybrid methodology proposed can most fruitfully be used in conjunction with sophisticated EKF/ SEIF based SLAM algorithms, wherein the SEIF/ EKF solve the Bayesian sub-problem of the hybrid formulation, in order to scale to really large, dense environments such as cities, ocean floors and planetary terrain.
The RBPF based SLAM algorithms, which falls under the broader category of trajectory based SLAM, on the other hand keep track of the whole trajectory of the robot which decorrelates the observations of the various map components. These methods have become very popular over the past few years and efficient techniques have been developed to form sparse landmark based maps as well as dense occupancy grid maps [7]–[10]. Another trajectory based method, called Consistent Pose Estimation (CPE) [11] relies on maintaining a graph on the poses at which various scans of the map were made and then, optimizing the inter-node distances such that the likelihood of the observed data is maximized given the statistics of the observation process [12]–[14]. However, these trajectory based methods keep track of the whole vehicle trajectory and thus, their state space grows unbounded over time. The CPE methods also require multiple passes over the same data to solve the pose optimization problem and thus, are essentially an offline batch processing method. For long term SLAM in large environments, it is still necessary to truncate the data at some finite time in the past [3] due to the constraints on the memory requirements, which may lead to the loss of consistency. In contrast, the method presented here is purely recursive and does not require the belief to be maintained over the entire vehicle pose history in order

Suman Chakravorty is an Assistant Professor, Aerospace Engineering, Texas A&M University, College Station, schakrav@aero.tamu.edu
R. Saha is a Graduate Research Assistant in the Aerospace Engineering Department at TAMU

978-1-4244-2794-9/09/$25.00 ©2009 IEEE

to decorrelate the map components, while being provably consistent.

It can be seen from the previous discussion that the main problem with the Bayesian formulation of the SLAM problem is the fact that computational complexity and consistency are at cross purposes, and alleviating one of the problems tends to worsen the other. In fact, starting with the important paper [1], the consistency of the SLAM algorithms, both RBPF and EKF/ SEIF based, has increasingly come under scrutiny in recent years [6], [15]–[17]. It was found that the RBPF based filters tend to lose consistency as time increases because of their inability to forget the past [16], [18], and the EKF based methods need stabilizing noise for consistency [15]. Also, neglecting the weak off-diagonal elements in SEIF based methods can lead to inconsistent results [6]. The method proposed in this paper first localizes with respect to a sparse set of landmarks in the map using a Bayes filter such as an EKF to obtain a belief on the pose of the robot, and then, maps the rest of the environment, based on the belief, using a frequentist approach. The frequentist part of the algorithm has complexity linear in the map and is provably consistent given that the Bayesian part of the problem is consistent. The complexity of the Bayesian part of the formulation can be kept under control owing to the sparseness of the set of landmarks/ features and using suitable sophisticated feature-based SLAM methods [19]. Further, the frequentist part of the formulation is immune to the data association problem, while the Bayesian part can be expected to be robust owing to the sparseness of the features/ landmarks and thus, the hybrid formulation is robust to the data association problem.

A valid question at this point is: why use the hybrid method? Why not localize using the EKF/ SEIF, and map the rest of the environment using a Bayesian method such as the occupancy grid OG method [20], as proposed in the DenseSLAM approach [21]. The answer to this question is that it may be impossible to maintain consistency in the Bayesian approach without maintaining correlation between map components, even under the "first localize - then map" philosophy adopted here! A simple counterexample is provided at the end of Section 2.4 to prove this. The other methods close to the work presented here, in philosophy, is [22], and other Expectation-Maximization (EM) based methods [23]. In these papers, a frequentist approach, the Baum-Welch algorithm [24], is used to find the ML estimate of a sparse set of landmarks and the vehicle trajectory, and then, this ML estimate is used to construct an OG map of the environment. However, the method used is an offline batch processing algorithm. The frequentist method proposed here is a recursive stochastic approximation algorithm and as such, different from the Baum Welch algorithm. In fact, it is well known in the Hidden Markov Model (HMM) literature that the recursive joint state-parameter estimation problem in HMMs cannot be robustly solved by modifying the B-W algorithm [25] (in the SLAM problem, the map is considered to be the unknown parameter). In fact, the use of frequentist estimators based on recursive ML (RML) or recursive least squares (RLS) is standard practice in the Hidden Markov Model literature [25], and an application of this methodology in the SLAM context is made in the reference [18]. These methods usually need to evaluate the filter derivative which is an $O(N^2)$ operation, where N is the number of particles used to represent the pdf of the state. This is usually impractical and only through recent advances in the particle filtering community [17], the above operation can now be done with $O(NlogN)$ complexity. In contrast, the method presented here does not require the filter derivative, and if the Bayesian part of the hybrid formulation is implemented using a particle filter, the complexity of the frequentist algorithms is $O(N)$, where N is the number of particles used to represent the robot pose pdf. This is accomplished by using (i) a probabilistic description of the map as the parameter in the methodology, instead of the deterministic description common in general joint state-parameter estimation algorithms [17], [18], [25], (ii) utilizing the "first localize -then map" philosophy, and (iii) by exploiting the structure of the resulting problem to intuitively define a frequentist estimator that is provably consistent.

An earlier version of this paper was presented at ACC 2008 [26]. The current paper extends the method to dense maps and solves the data association problem in such maps. The rest of the document is organized as follows. In section 2, we present the hybrid approach to the SLAM problem. In particular, in section 2.1, we present a frequentist alternative to Elfes' OG method to mapping with known poses. In section 2.2, we generalize the method to mapping with uncertain poses wherein the robot pose is specified by a time varying probability density function (pdf). In section 2.3, we present the hybrid method to solve the SLAM problem. In section 3, we present simulation experiments wherein large environments with multiple cycles are mapped using the hybrid methodology.

II. The Hybrid Bayesian/ Frequentist Methodology

A. Frequentist Mapping

Consider a single autonomous agent and let its state be denoted by the variable s (also sometimes called the robotic pose), and let the state of the environment be denoted by the variable $Q = \{q_1, \cdots, q_M\}$, where q_k are components of the environment (for instance, these would be the individual grid cells in a grid cell decomposition of the environment). The state and the environment are assumed to be discrete-valued random variables. The environment is assumed to be stationary and uncorrelated, i.e.,

$$p^*(Q) = \prod_{i=1}^{M} p^*(q_i), \qquad (1)$$

where $p^*(q_i)$ represents the "true" probability that the component q_i can take one of D possible values. For instance, in the OG approach, this would correspond to the "true" probability that a particular grid in the map is occupied or not. It can be anticipated that an overwhelmingly large part of most environments can be modeled in this fashion. In fact, any deterministic environment trivially satisfies the above assumptions. The probability of observing the i^{th} environmental component in the state $\hat{q}_i$, where $\hat{q}_i$ can take one of D values, and given that it is observed from the pose s, is given by:

$$p(\hat{q}_i/s) = \sum_{q_1,\cdots,q_N} p(\hat{q}_i/q_1,\cdots,q_N,s)p^*(q_1)\cdots p^*(q_N). \quad (2)$$

The above equation can be rewritten as:

$$p(\hat{q}_i/s) = \sum_{q_i} p^*(\hat{q}_i/q_i,s)p^*(q_i), \quad (3)$$

$$p^*(\hat{q}_i/q_i,s) = \sum_{q_1,.,q_{i-1},q_{i+1},.,q_N} p(\hat{q}_i/q_1,.,q_N,s)p^*(q_1)..p^*(q_{i-1})p^*(q_{i+1})..p^*(q_N). \quad (4)$$

The above may be compactly written in matrix form as the equation

$$\hat{P}_i(s) = A_i^*(s)P_i^*, \quad (5)$$

where the vector $\hat{P}_i(s)$ stacks the observation probabilities $p(\hat{q}_i/s)$, and the matrix $A_i^*(s)$ is the true observation model of the i^{th} component when observed from pose s. **The above equation is the fundamental equation for the frequentist approach and provides an avenue for estimating the true environmental probabilities** P_i^*. Suppose we make repeated observations of the i^{th} component from pose s. We could count the number of times that we observe the i^{th} component in its various states, and form a consistent estimate of the observation probability vector $\hat{P}_i(s)$ by averaging, i.e.,

$$\hat{P}_i(s) = E_z[1(\hat{q}_i/s,z)] \equiv E_z[c_i(s,z)] = \lim_N \frac{1}{N}\sum_{t=1}^{N} 1(\hat{q}_{i,t}/s,z_t). \quad (6)$$

In the above, given an observation z, the observation vector $c_i(s,z) = [1(\hat{q}_i/s,z)]$ ($1(.)$ denotes the indicator function) enters a one into the $\hat{q}_i$ component, and zero in every other component at time t (for instance, in the occupancy grid representation it will enter a 1 into the "occupied" entry if the grid cell is observed to be occupied, or a 1 into the "empty" entry otherwise). The above equation is correct due to the Law of Large Numbers. Then, using the knowledge of $A_i^*(s)$, we can obtain the true environmental probabilities P_i^* as

$$P_i^* = A_i^*(s)^{-1}\hat{P}_i(s). \quad (7)$$

Next, we may relax the assumption that the observations are made from the pose s and have that the observations are made from the time varying poses $\{s_t\}$, with true observation models $A_i^*(s_t)$. Again, if we keep track of the relative frequencies of observations of the i^{th} component in its various different states, then the estimate of the true probabilities P_i^* can be recovered asymptotically using a time averaged observation model as follows:

$$P_i^* = \bar{A}_i^{-1}\hat{P}_i, \quad (8)$$

$$\hat{P}_i = \frac{1}{N}\sum_{t=1}^{N} c_i(s_t,z_t), \quad (9)$$

$$\bar{A}_i = \frac{1}{N}\sum_{t=1}^{N} A_i^*(s_t). \quad (10)$$

If we interpret the frequency of seeing the i^{th} map component in its $\hat{q}_i$ level during the course of the mapping experiment as a probability, and if we interpret the frequency of the robot being in a state s as a probability, then it follows using the simple rules of conditional probability that:

$$p(\hat{q}_i) = \sum_{q_i,s} p^*(\hat{q}_i/q_i,s)p^*(q_i)p(s) = \sum_{q_i}[\sum_{s} p^*(\hat{q}_i/q_i,s)p(s)]p^*(q_i). \quad (11)$$

Provided that the state s_t converges to some stationary distribution, the left hand side $p(\hat{q}_i)$ in the above equation is given by Eq. 9, and the matrix $[\sum_s p^*(\hat{q}_i/q_i,s)p(s)]$ is given by Eq. 10, and hence, the estimation equations for the time varying pose case follow. The true environmental probabilities can then be recovered recursively using the following estimator if $A_i^*(s_t)$ is positive definite (which is true under mild conditions).

Estimator E1:

$$P_{i,t} = \Pi_{\mathcal{P}}\{P_{i,t-1} + \gamma_t(c_i(s_t,z_t) - A_i^*(s_t)P_{i,t-1}\}, \quad (12)$$

where $\mathcal{P}$ represents the space of probability vectors in $\Re^D$, and $\Pi_{\mathcal{P}}(.)$ denotes a projection onto this compact set. The sequence $\{\gamma_t\}$ is usually of the form $at^{-\alpha}, \alpha < 1$, where a and α are design parameters, we have usually used $\alpha = -1$ adn $a = 0.1$ in our simulations. However, there still remains the problem of using the "true" observation models $A_i^*(s)$ in order to form the estimates. We have an estimate $P_i(t)$ of the map probabilities for the different components and these estimates are used in Eqs. 4 -5 to form the observation models $A_i(s)$ as an approximation of the true observation models. These models can be inferred from the model of the particular type of sensor used for sensing the environment [20].

B. Frequentist Mapping with Uncertain Robotic Poses

In this section, we relax the assumption that the pose of the robot is known perfectly. Instead, we assume that we are given a belief, i.e., a probability distribution, on the possible poses of the robot. Given the belief on the pose of the robot, $b_t(s)$ at time t, and a reading z_t of the environment, the frequentist mapping method is now used to map the (dense) environment Q. However, it is immediately apparent that there is an inherent "data association" problem associated with the mapping problem in this scenario. The observation, $\hat{q}_i$, of an environmental component q_i is no longer certain, since it varies with the pose of the robot. Consider the simple situation illustrated in Fig. 1. The map component q_1, given reading z_2, is empty or occupied depending on whether the robot is at pose s_1 or s_2 respectively. Thus, given the uncertainty in the pose of the robot $b(s)$ and the reading of the environment z, the observation of the i^{th} component of the environment $\hat{q}_i$

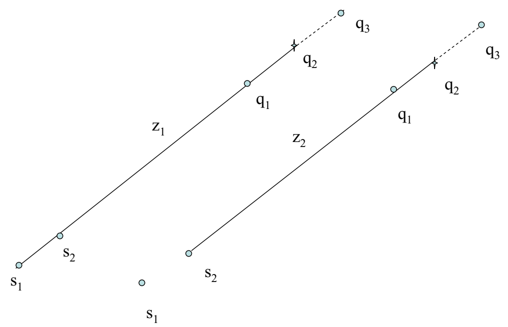

Fig. 1. The problem of data association

is given by the probability vector (derived using the rules of conditional probability, and Bayes rule)

$$c_i^*(b(s),z) \equiv [p(\hat{q}_i/b,z)] = \sum_s [1(\hat{q}_i/s,z)] \frac{p^*(z/s)b(s)}{p^*(z/b)}, \quad (13)$$

$$p^*(z/s) = \sum_{q_1,\cdots,q_N} p(z/s,q_1,\cdots q_N)p^*(q_1)\cdots p^*(q_N), \quad (14)$$

where $p^*(z/b) = \sum_s p^*(z/s)b(s)$ is the factor used to normalize $c_i(.)$ and $p^*(z/s)$ is the true likelihood of the observation z given that it is made from pose s. In order to derive the above expression, note that using the thorem of total probability, $p(\hat{q}_i/b,z) = \sum_s p(\hat{q}_i/s,b,z)p(s/z,b)$. We can expand the term $p(s/z,b)$ using Bayes rule which gives us $p(s/z,b) = \frac{p^*(z/s)b(s)}{p^*(z/b)}$, and using the fact that $p(z/s,b) = p(z/s)$, Eq. (13) above follows.

As in the perfect pose information case, averaging over all observations z (which can be formed by a time average due to the Law of Large Numbers), allows us to estimate the probability of observing state $\hat{q}_i$ given the belief state $b(s)$, i.e.,

$$p(\hat{q}_i/b) = E_z[c_i^*(b,z)] \approx \frac{1}{N}\sum_{t=1}^{N} c_i^*(b,z_t). \quad (15)$$

Note that the above probabilistic description of the observation solves the "data association" problem: we are no longer certain if the observed value of the i^{th} map component is in its k^{th} level, instead we associate a probability with this observation. The probability of observing the map component q_i at level $\hat{q}_i$, given the belief on the pose $b(s)$ is also given by

$$p(\hat{q}_i/b) = \sum_s b(s) \sum_{q1\cdots,q_N} p(\hat{q}_i/q1,\cdots,q_N,s) \\ p^*(q_1)\cdots p^*(q_N), \quad (16)$$

which can be written in compact matrix form as follows:

$$\hat{P}_i(b) = A_i^*(b)P_i^*, \quad (17)$$

where $\hat{P}_i(b) = [p(\hat{q}_i/b)]$, and

$$A_i^*(b) = \sum_s A_i^*(s)b(s). \quad (18)$$

Note here that this equation is exactly analogous to the frequentist mapping equation 5, wherein the exact pose knowledge s has been replaced by the belief on the pose of the robot $b(s)$. The observation model $A_i^*(s)$ is replaced by the averaged observation model with the averaging being done with respect to the belief on the pose of the robot. Thus, similar to the case with perfect pose information, if we were to remain in the belief state $b(s)$ and make repeated observations of the i^{th} component of the environment, we would be able to recover the left hand side of the above Eq. 17, $\hat{P}_i(b)$, by averaging the (probabilistic) observations of the i^{th} component, $c_i(b,z_t)$ (cf. Eq. 15). Hence, the true environmental probabilities may be recovered asymptotically by inverting Eq. 17. Generalizing the situation to the case when we have a time-varying belief on the pose of the robot, $b_t(s)$, the true environmental probabilities can be estimated recursively using the following analog of frequentist estimator E1 .

Estimator E2

$$P_{i,t} = \Pi_{\mathcal{P}}\{P_{i,t-1} + \gamma_t(c_i^*(b_t,z_t) - A_i^*(b_t)P_{i,t-1})\}, \quad (19)$$

As in the pure mapping case, the estimator is actually run by using the current estimate of the true observation models/ observation likelihood. In other words, the above algorithm is run using $c_i(b_t,z_t,P_t)$ and $A_i(b_t,P_t)$, where the current estimate of the map probabilities P_t is used, instead of the true map probabilities P^*, in Eq. (13) to form $c_i(b_t,z_t,P_t)$, and in Eqs. (16)-(17) to form $A_i(b_t,P_t)$.

C. Hybrid Bayesian/ Frequentist SLAM

At this point, we formulate a hybrid methodology to generalize the frequentist mapping methodology to the simultaneous localization and mapping (SLAM) case. It is assumed that the system localizes itself with respect to a sparse, well separated set of features/ landmarks $\Theta = \{\theta_1,\cdots,\theta_K\}$. Then, the belief (or probability distribution) over the pose-features pair is formed recursively using a Bayes filter (such as a Kalman filter in the Linear Gaussian case):

$$b_t(s,\Theta) = p(z_t^\theta/\Theta,s)\sum_{s'} p(s/s',u_{t-1})b_{t-1}(s',\Theta), \quad (20)$$

where z_t^θ represents the noise corrupted observation of the landmarks Θ at time t, and u_{t-1} denotes the control acting on the system at time $t-1$. The identification and recognition of these features and landmarks in an autonomous fashion is a challenging problem in itself, but can be solved using suitable feature-based SLAM algorithms [2], [4], [5]. Given the joint distribution of the pose-landmark pair, the belief on the pose of the vehicle is formed by marginalizing the dependence on the landmarks, and is output to the frequentist part of the mapping algorithm. The frequentist part of the method, i.e., Estimator E2, is now used to map the rest of the (dense) environment using the belief output from the Bayesian part of the methodology. Thus, the hybrid methodology can be represented as the following algorithm:

Given a reading of the map z_t, and given that the Bayesian part of the hybrid problem formulation is solved using a

Hybrid SLAM

Given $b_0(s)$, initial map occupancy probabilities $P_i(0)$, and reading of environment z_1, $t = 0$. ***Do** till convergence*

- **Bayesian:** Extract the readings of the landmarks, z_t^{θ}, from the raw sensor readings z_t, and form the belief on the state of the robot, $b_t(s)$, using Eq. (20) and marginalizing over the landmarks.
- **Frequentist:** Take the rest of the data z_t^Q. Isolate each component of the (dense) map Q that is observed given z_t^Q (cf. Eq. 13), and the belief $b_t(s)$ from the step above. Update the map probabilities of each component using the belief $b_t(s)$ in the frequentist estimator E2 (cf Eq. 19).

End

particle filter, the computational complexity of the frequentist estimator in updating any map component is $O(N)$ where N is the number of particles used to represent the belief state. The Bayesian part of the formulation inherits the computational complexity of whatever method is used to solve that part of the problem. Also, note that each map component is updated completely independent of the others and hence, the method has complexity linear in the map components, i.e., $O(M)$, where M is the number of map components. Contrast this with the $O(M^2)$ complexity of the Kalman filter based approach. In order to make this clear, suppose that there are $M+N$ total components in a map. At the basic algorithmic level, in the Kalman filter based approach, the computational complexity is $O(N+M)^2$. In the Hybrid formulation, suppose that N is the number of features that is used to localize the robot and M is the rest of the map. Then, at the basic algorithm level, the computational complexity of the hybrid approach is $O(N^2) + O(M)$. Thus, if $N << M$, then the hybrid formulation possesses orders of magnitude better computational benefits compared to Bayesian methods such as the EKF. Moreover, due to the sparseness of the landmark/ features, the data association problem for the Bayesian sub-problem is significantly simpler. In conjunction with the robustness of the frequentist estimator to the data association problem, this leads to significantly improved robustness of the hybrid formulation to the data association problem.

Due to the paucity of space, we cannot provide the proof of the above method in this paper. Also, that the frequentist part of the hybrid technique is necessary can be proved through a simple counterexample which is left out as well due to the space constraints here. The interested reader may find these items in the extended technical report on the first author's webpage.

III. EXPERIMENTS

The hybrid SLAM methodology has been applied to multiple simulated large scale maps and the results bear out the theoretical guarantees that the method possesses. The robot model used in the simulations was that of a differential drive vehicle [27]. The dimensions of the robot were as follows: a

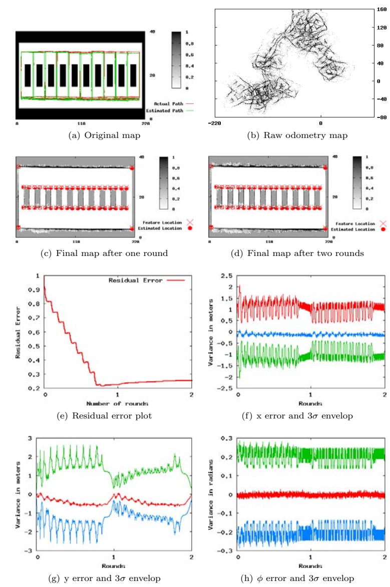

Fig. 2. Mapping using noisy 2-D sensor (Map 1)

wheel radius of 25 cm, and a width of 50 cm. Experiments were performed for two different kinds of sensors: a) a noisy 2-D sensor with both range and bearing errors (such as a sonar) with σ_r = 0.2 m and σ_θ = 0.6 deg, and b) an accurate 1-D range sensor such as a SICK laser range sensor with σ_r = 0.01 m and σ_θ = 0.05 deg. The process noise in the wheel encoders was σ_u = 0.5 rad/s. The average robot wheel speed was 5 rad/sec and the integration time step for the EKF was 0.5 sec. Here, we give representative results for a large map of dimensions 220m x 40m (Fig. 2(a)). The corners of the corridors in the map were the features used for localization and were assumed to be reliably identified. Due to their sparseness, they are at least 20m apart, we do not consider the data association problem. The robot localizes itself with respect to these features using an EKF and then, based on the belief on the robot pose output by the EKF, it maps the rest of the environment using the frequentist estimator E2. The robot loops every corridor from left to right before continuing on to the next corridor and repeats this process twice, i.e., makes two rounds or laps of the map. The length of the robot run was approximately 2.5 Km. The results of the mapping runs are shown in Figs. 2 and 2. For Map 1 (the 220m x 40m map), the results for the run with the noisy sensor are shown in Fig. 2 and for the accurate laser sensor in Fig. 3. In Figs. 2/3, Subfigure (a) shows the original map along with the true robot path overlaid with the estimated robot path, Subfigure (b) shows the raw odometry data, Subfigures (c) and (d) shows the estimated map, along with the features and their estimates, after the completion of one and two rounds/ laps respectively, Subfigure (e) shows the

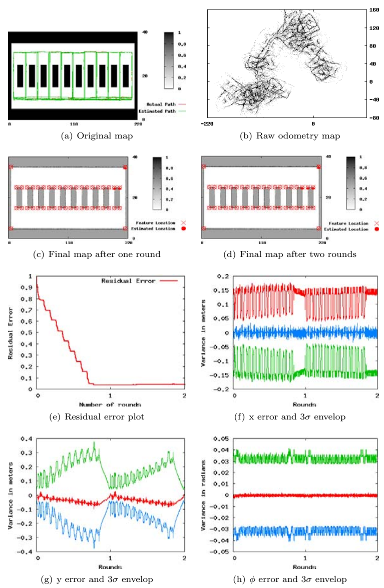

(a) Original map (b) Raw odometry map (c) Final map after one round (d) Final map after two rounds (e) Residual error plot (f) x error and 3σ envelop (g) y error and 3σ envelop (h) ϕ error and 3σ envelop

Fig. 3. Mapping using accurate laser sensor(Map 1)

total error in the map, which is the fraction of grids that have not converged to their true occupancy values, as a function of the number of laps the robot makes, Subfigures (f)-(h) show the error in the x, y, and θ co-ordinates of the robot along with their associated 3σ bounds. The method has been tested on several other large maps with multiple loops as well but we cannot present all these results her due to the space constraints. These figures give us an idea as to how well the algorithm is performing and also give us valuable practical insight into the algorithm. One of the reasons we chose such an example is because of the well-known challenge maps with multiple cycles pose to SLAM algorithms which is evidenced from the raw odometry plots (Subfigure (b) in the plots). These plots show that the scale and difficulty level of these mapping problems is on par with known datasets such as the ones on OPENSlam.org, except that our data was simulated instead of being from a real experiment. The algorithm had no problems in closing large loops as the ones shown here and we did not have to make any heuristic corrections when such a loop was closed. In fact, the size of the map, or the number of cycles in it, is really not a problem for this algorithm as long as the EKF remains consistent. However, if the EKF loses consistency then the guarantees of consistency for the frequentist part of the hybrid formulation are no longer valid. Thus, sophisticated Kalman/ Information filter based methods for feature-based SLAM can play an important role in ensuring the consistency of the hybrid formulation for maps on a much larger scale such as cities, planetary terrains etc., where the order of the distances are in hundreds or thousands of kilometers, and consequently, the number of features/ landmarks increases by several orders of magnitude when compared to the maps shown here. It can also be seen from the total map error plots (Subfigure (e) in the plots) that the mapping algorithm converges exponentially. In fact, it is our conjecture that this can be rigorously established, and is supported by the experimental evidence. Hence, we may conclude that the preliminary results show sufficient evidence of the efficacy/ applicability of the methodology proposed here.

REFERENCES

[1] M. W. M. Dissanayake, P. Newman, S. Clark, and H. F. Durrant-Whyte, "A solution to the simultaneous localization and map building (slam) problem," *IEEE Transactions on Robotics and Automation*, vol. 17, pp. 229–241, 2001.

[2] J. Leonard and H. Durrant-Whyte, *Directed Sonar Sensing for Mobile Robot Navigation*. Boston, MA: Kluwer academic, 1992.

[3] T.Bailey and H. Durrant-Whyte, "Simultaneous localization and mapping (slam): Part ii," *IEEE Robotics and Automation Magazine*, pp. 108–117, Sep. 2006.

[4] S. Thrun, Y. Liu, D. Koeller, A. Y. Ng, S. Ghahramani, and H. Durrant-Whyte, "Simultaneous localization and mapping with sparse extended information filters," *International Journal of Robotics Research*, vol. 23, pp. 693–716, 2004.

[5] R. Eustice, H. Singh, and J. Leonard, "Exactly sparse delayed state filters," in *Proc. IEEE Int. Conf. Rob. Aut.*, 2005.

[6] M. Walter, R. Eustice, and J. Leonard, "A provably consistent method for imposing sparsity in feature-based slam information filters," in *S. Thrun, R. Brooks, H. Durrant-Whyte (Eds.), Robotics Research, STAR 28*, 2007, pp. 214–234.

[7] M. M. et. al., "Fastslam 2.0: An improved particle filtering technique for simultaneous localization and mapping that provably converges," in *Proc. of IJCAI*, 2003.

[8] K. Murphy, "Bayesian map learning in dynamic environments," in *Advances in Neural Information Processing Systems (NIPS)*, 1999.

[9] A. Eliazar and R. Parr, "Dp-slam: Fast, robust simultaneous localization and mapping without predetermined landmarks," in *Proceedings of the Eighth International Joint Conference on Artificial Intelligence (IJCAI 03*, 2003.

[10] G. Grisetti, C. Stachnis, and W. Burgard, "Improved techniques for grid mapping with rao-blackwellized particle filters," *IEEE Trans. Robotics*, vol. 23, pp. 34–46, 2007.

[11] H. Durrant-Whyte and T. Bailey, "Simultaneous localization and mapping (slam): Part i," *IEEE Robotics and Automation Magazine*, pp. 99–108, June 2006.

[12] F. Lu and E. Milios, "Globally consistent range scan alignment for environment mapping," *Autonomous Robots*, vol. 18, pp. 249–275, 1997.

[13] J. S. Gutmann and K. Konolige, "Incremental mapping of large cyclic environments," in *Proc. IEEE Int. Symp. on Comp. Intell. in Rob. and Aut. (CIRA)*, 2000.

[14] S. Thrun and M. Montemerlo, "The GraphSLAM algorithm with applications to large-scale mapping of urban structures," *International Journal on Robotics Research*, vol. 25, no. 5/6, pp. 403–430, 2005.

[15] T. Bailey, J. Nieto, , J. Guivant, M. Stevens, and E. Nebot, "Consistency of the ekf-slam algorithms," in *Proceedings of the IEEE Int. Conf. Rob. automat.*, 2006.

[16] T. Bailey, J. Nieto, and E. Nebot, "Consistency of the fastslam algorithm," in *Proceedings of the IEEE Int. Conf. Rob. automat.*, 2006.

[17] M. Klaas, N. DeFreitas, and A. Doucet, "Towards practical n^2 monte carlo: the marginal filter," in *Uncertainty in Artificial Intelligence*, 2005.

[18] R. Martinez-Cantin, N. deFreitas, and J. A. Castellanos, "Analysis of particle methods for simultaneous robot localization and mapping and a new algorithm: Marginal-slam," in *Proceedings of the IEEE Int. Conf. Rob. automat.*, 2007.

[19] J. Guivant and E. Nebot, "Optimization of the simultaneous localization and map-building algorithm for real-time implementation," *IEEE Tr. Rob. Aut.*, vol. 17, pp. 242–257, 2001.

[20] A. Elfes, "Using occupancy grids for mobile robot perception and navigation," *IEEE Computer*, vol. 22, pp. 46–57, 1989.

[21] J. N. et. al., "The hybrid metric maps (hymm)," in *IEEE Int. Conf. Rob. Automat.*, 2004.

[22] S. Thrun, D. Fox, and W.Burgard, "A probabilisitic approach to concurrent mapping and localization for mobile robots," *Mach. Learn. Aut. Robots*, vol. 31, pp. 29–53, 1998.

[23] H. Shatkay and L. P. Kaelbling, "Learning topological maps with local odometric information," in *Proc. IJCAI-97*, 1997.

[24] L. R. Rabiner, "A tutorial on hidden markov models and selected applications in speech recognition," *Proceedings of the IEEE*, vol. 77, pp. 257–286, 1988.

[25] Y. Ephraim and N. Merhav, "Hidden markov processes," *IEEE Transactions Inf. Theory*, vol. 48, pp. 1518–1569, 2002.

[26] S. Chakravorty, "A hybrid approach to the simultaneous planning localization and mapping problem," in *Proceedings of the American Control Conference*, 2008.

[27] S. M. LaValle, *Planning Algorithms*. Cambridge, UK: Cambridge University Press, 2005.