A Rao-Blackwellized Particle Filter for Topological Mapping

Ananth Ranganathan and Frank Dellaert
{ananth,dellaert}@cc.gatech.edu
College of Computing
Georgia Institute of Technology, Atlanta, GA

***Abstract*— We present a particle filtering algorithm to construct topological maps of an uninstrument environment. The algorithm presented here constructs the posterior on the space of all possible topologies given measurements, and is based on our previous work on a Bayesian inference framework for topological maps [21]. Constructing the posterior solves the perceptual aliasing problem in a general, robust manner. The use of a Rao-Blackwellized Particle Filter (RBPF) for this purpose makes the inference in the space of topologies incremental and run in real-time. The RBPF maintains the joint posterior on topological maps and locations of landmarks. We demonstrate that, using the landmark locations thus obtained, the global metric map can be obtained from the topological map generated by our algorithm through a simple post-processing step. A data-driven proposal is provided to overcome the degeneracy problem inherent in particle filters. The use of a Dirichlet process prior on landmark labels is also a novel aspect of this work. We use laser range scan and odometry measurements to present experimental results on a robot.**

I. Introduction

The last decade of research in the field of robotic mapping has seen the emergence of two major paradigms - metric maps and topological maps. Metric maps [18][17], which have enjoyed more popularity, provide a fine-grained, scale-consistent representation of the environment with the added benefit of being easy to use for navigation. However, they are hard to build accurately due to their fine-grained structure. Topological maps [23][4], on the other hand, are a graph-based description of the environment that offer an orthogonal set of benefits and drawbacks in the sense that they are a coarse representation but not directly amenable for use in navigation. For a detailed discussion comparing metric and topological maps, see [27]. In this work, we use the standard definition of a topological map, i.e a set of landmark nodes connected by edges that denote traversibility [13].

A major problem in topological mapping is perceptual aliasing in the environment, whereby many distinct landmarks appear to be similar to the robot's sensors. Thus, the robot has trouble labeling the landmark correctly, and consequently, in inferring the correct topology. Previous work in this area has attempted to solve the perceptual aliasing problem through the use of Hidden Markov Models (HMMs) [25][8], multiple hypothesis tracking using POMDPs [29], maintaining a tree of all consistent topologies at each step [24], and clustering of measurements from the landmark locations [12]. Other approaches use local appearance at a landmark to perform place recognition [6][26]. However, many of these approaches are brittle and prone to silent failure in difficult environments.

Recently, we have proposed a Bayesian framework for inferring the posterior distribution on the space of topologies [21][22]. This technique, called Probabilistic Topological Maps (PTMs) provides a general, robust solution to the problem of perceptual aliasing. PTMs are generated through inference in the space of topologies, which is equivalent to the space of all possible set partitions of the landmark measurements. This latter space is combinatorial in nature and the number of possible topologies becomes intractably large even for a modest number of distinct landmarks.

As our primary contribution in this paper we propose a particle filtering algorithm for topological mapping based on the PTM framework using a laser scanner as sensor. The use of particle filters makes our algorithm incremental, as opposed to our previous work on PTMs that used Markov Chain Monte Carlo (MCMC) [21], which is inherently a batch algorithm. While the use of importance sampling in a highly combinatorial space may seem suspect, it is justified by the fact that the posterior in the space of topologies is highly peaked; only a handful of topologies get a non-zero probability mass. Additionally, the use of laser range scanners in the PTM framework is also a novel aspect of this work. Range scans are used to construct local map patches around landmark locations that the robot visits. These map patches are subsequently matched using scan matching techniques to provide a likelihood of their being from the same physical location.

We present a data-driven proposal distribution that makes use of odometry measurements to overcome the samples degeneracy problem in the particle filter and encourage fast convergence. A novel prior distribution on the type of landmark that the robot expects to see next is also provided. This prior takes the form of a Dirichlet process and encodes intuitive characteristics of the problem domain.

Our algorithm uses a Rao-Blackwellized Particle Filter (RBPF) [19][17] to maintain a joint posterior on the landmark locations and topologies. Hence, the posterior distribution on the metric locations of the landmarks can be obtained as a side effect of inferring the topology. A Lu-Milios style smoothing operation that incorporates topological constraints [9], performed as a post-processing step after the addition of a landmark to the PTM, can be used to produce a metric

Fig. 1. Scan measurements, obtained by concatenating scans from around landmark locations, used by the RBPF algorithm.

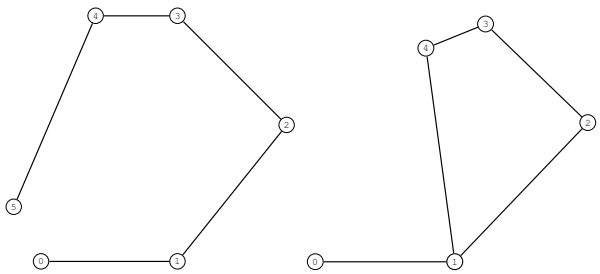

(a) (b)

Fig. 2. Two topologies with 6 observations, each corresponding to set partitions (a) with six landmarks $(\{0\},\{1\},\{2\},\{3\},\{4\},\{5\})$ and (b) with five landmarks$(\{0\},\{1,5\},\{2\},\{3\},\{4\})$, illustrate the equivalence between topologies and set partitions.

map of reasonable accuracy. Combined with the topological map, this metric information makes navigation using PTMs simple. Also, since a PTM is a posterior on topologies, this step gives us a posterior on hybrid metric-topological maps of the environment. As we have an estimate of the correctness of any map in the posterior, navigation can be done in a robust manner using an ensemble of maps.

In this work, we do not deal with the problem of landmark detection, which is largely an orthogonal issue to the main problem of inferring topologies. Any landmark detection scheme, such as the ones in [15][12], can be used with our algorithm. The only assumption made is that the output of the landmark detection operator does not have false negatives, i.e landmarks do not get skipped over without the robot sensing them. However, as our algorithm is capable of dealing with false positives from the landmark detector, any of the above mentioned schemes can be tuned to avoid false negatives while giving a fair number of false positives. Thus our assumption in this regard does not restrict the use of our algorithm in any way. We present experiments performed on an ATRV-Mini robot to validate our algorithm that make use of manually selected landmark locations.

II. The Space of Topologies

Our aim is to compute the posterior over topologies given measurements, $P(T^n|z^n)$, where T^n is a topology constructed from n landmarks observed by the robot in its run so far, and z^n is the set of measurements upto the nth landmark. In our case, the measurements consist of odometry and laser range scans, so that $z^n = \left\{s^n, o^{n-1}\right\}$, where $s^n = \{s_1, s_2, \ldots, s_n\}$ is the set of range measurements and $o^{n-1} = \{o_1, o_2, \ldots, o_{n-1}\}$ is the set of odometry measurements between landmarks. The scan measurements, used in one of our experiments, are given in Figure 1. Before we can perform inference in the space of topologies, we need to understand the nature of this space.

A. Topologies as set partitions

The key idea behind inference in the space of topologies is the equivalence between topologies of an environment and set partitions of landmark measurements. A set partition on the measurements groups them into a set of equivalence classes in which each set is associated with a distinct landmark. When all the measurements from the same physical landmark are clustered together, this naturally defines a partition on the set of measurements. An example of the one-to-one correspondence between topologies and set partitions is shown in Figure 2. Let us assume, as before, that the robot has observed n landmarks so far. If the number of distinct landmarks out of these is m ($m \leq n$), then a topology T^n can be represented as the set partition of the set measurement z, $T^n = \{S_j \mid j \in [1, m]\}$, where each S_j is a set of measurements such that $S_{j1} \cap S_{j2} = \phi$ $\forall j1, j2 \in [1, m]$, $j1 \neq j2$ and $\bigcup_{j=1}^{m} S_j = z^n$. The set S_j contains the measurements corresponding to the jth physical landmark in the environment.

It can be seen that a topology is nothing but an assignment of measurements to sets in the partition. This results in the above mentioned isomorphism between topologies and set partitions. The number of possible topologies for a given environment is thus equal to the number of set partitions of the set of measurements. This number is called the *Bell number* [20], and grows hyper-exponentially with the number of measurements.

If we associate a label with each landmark, we can also represent the topology T^n by a label sequence $L^n = L_{1:n}$, where L_i is the label of the ith landmark. Further the number of unique labels in this sequence is equal to the number of sets in the set partition corresponding to the topology T^n, i.e. m. The posterior on topologies that we seek can then be written as $p\left(L^n|z^n\right)$.

III. Particle Filtering for Topological Mapping

Since the space of topologies is discrete and combinatorial in size, it is not possible to compute the posterior in analytical form. Probabilistic Topological Maps (PTMs), which we introduced in earlier work [21], overcome this problem by

Symbol	*Meaning*
n	Total number of landmarks observed
m	Number of distinct landmarks observed
o^{n-1}	The n odometry measurements
s^n	Range scan measurements around the n landmarks
z^n	Combined set of measurements $z^n = \left\{ s^n, o^{n-1} \right\}$
L^n	Topology T^n represented as a label sequence
X^n	Landmark locations for the topology L^n
$\alpha_n\left(X^n\right)$	Analytic distribution on the landmark locations

TABLE I

NOTATION USED IN THE EXPLANATION OF THE ALGORITHM

maintaining a histogrammed, sample-based approximation to the posterior distribution over topologies given some measurements. In this work, we use the particle filtering framework to compute this sample-based approximation. A summary of the notation used in our exposition of the algorithm is given in Table I.

The posterior on topologies that we seek is represented as $p(L^n|s^n, o^{n-1})$. Applying Bayes law on the required posterior to obtain the measurement likelihood and prior, we get

$$p(L^n|s^n, o^{n-1}) \quad \propto \quad p(L^n|z^{n-1})p(s_n, o_{n-1}|L^n, z^{n-1}) \quad (1)$$

where the measurements upto the $(n-1)$th landmark have been represented as $z^{n-1} = \left\{ s^{n-1}, o^{n-2} \right\}$ and the likelihood of the measurements from the nth observed landmark is $p(s_n, o_{n-1}|L^n, z^{n-1})$. The prior $p(L^n|z^{n-1})$ can be further factorized to give an incremental prior on the label in the current (nth) time step -

$$p(L^n|z^{n-1}) \quad = \quad p(L_n|L^{n-1}, z^{n-1})p(L^{n-1}|z^{n-1}) \quad (2)$$

where $p(L_n|L^{n-1}, z^{n-1})$ is the prior (proposal) distribution for the label on the nth observed landmark and $p(L^{n-1}|z^{n-1})$ is the posterior from the previous step containing $n-1$ measurements. The prior gives a distribution on which of the distinct landmarks we are likely to see next, including the possibility of the next landmark being a previously unvisited one.

Since our algorithm is based on a particle filter, we represent the posterior by a set of weighted particles

$$S_n \quad = \quad \left\{ L^{n,(i)}, w_n^{(i)} \right\}_{i=1}^{N} \quad (3)$$

where $w_n^{(i)}$ is the weight on the ith particle. It can be seen that equations (1) and (2) together give a recursive formulation for the posterior on topologies that is amenable for performing particle filtering. An illustration of a set of samples from the particle filter is given in Figure 3.

The two components required to perform filtering are the proposal distribution and a method for computing the importance weights. These are explained in the following sections.

A. The Proposal Distribution

We use the predictive prior distribution on the current landmark label $p(L_n|L^{n-1}, z^{n-1})$, given in (2), as our proposal distribution. Using the sample notation of (3), the proposal distribution can be written as

$$L_n^{(i)} \quad \sim \quad p\left(L_n|L^{n-1,(i)}, z^{n-1}\right) \quad (4)$$

This is a discrete probability distribution on a vector of size $p+1$, where p is the number of distinct landmarks observed upto the $(n-1)$th step. The distribution (4) encodes our expectation of the robot revisiting one of the previously observed landmarks or visiting a completely new one. While a uniform distribution can be used for this purpose, it does not capture all the characteristics of the problem. For example, with a uniform distribution, the probability of a robot visiting a new landmark remains constant with time. However, we can reasonably expect the robot to visit fewer new landmarks as its run progresses. Similarly, landmarks that have been visited frequently in the past should be better candidates for revisitations. This is especially true for indoor environments where lobbies and corridor junctions are visited more frequently than other locations.

A distribution that models these problem characteristics well is the Dirichlet process prior [7][1]. The Dirichlet process is an extension of the standard Dirichlet distribution to infinite mixture models, i.e it also includes a probability of observing previously unobserved measurement classes, in our case landmarks. The prior on the nth landmark label using the Dirichlet process is given as

$$p\left(L_n^{(i)}|L^{n-1,(i)}\right) \quad = \quad \begin{cases} \frac{n(L_j^{(i)})}{n+c} & 1 \leq j \leq p \\ \frac{c}{n+c} & j = p+1 \end{cases} \quad (5)$$

where p is the number of distinct landmarks observed upto the $(n-1)$th step as before, and $n(L_j^{(i)})$ is the number of occurences of the label j in the label sequence corresponding to the topology.

The parameter c encodes our belief in the number of distinct landmarks in the environment. A large value of c increases the probability of observing a new landmark at every step and consequently, the number of distinct landmarks in the topology. Note that the probability of observing a new landmark decreases as n increases though it never goes down to zero. Also, the probability of revisiting a landmark is proportional to the number of times it has been visited before, given by $n(L)$.

B. Rao-Blackwellization in the Importance Weight Computation

Using the definition of the importance sampling weights, we see that the expression for the importance weights is the same as the measurement likelihood

$$w_{m,n}^{(i)} \quad = \quad \frac{\text{Target distribution}}{\text{Proposal distribution}} \quad (6)$$

$$\propto \quad p(s_n, o_{n-1}|L^{n,(i)}, z^{n-1}) \quad (7)$$

where we have used the target distribution from (1) and proposal from (2).

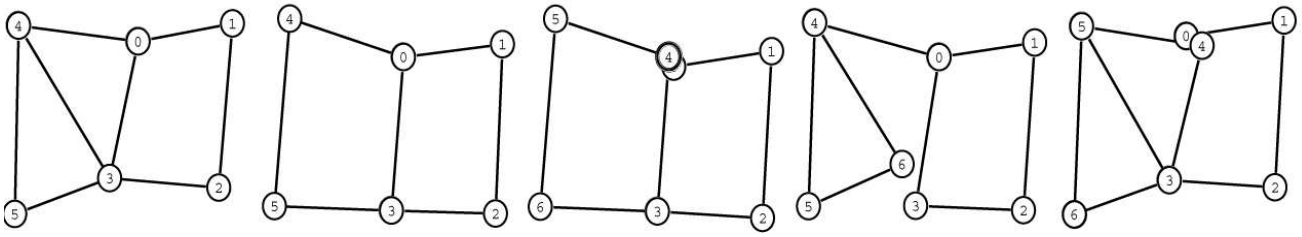

Fig. 3. Example of a set of samples from the space of topologies for an environment. Each sample is associated with a weight in the particle filter.

We introduce the landmark locations by marginalizing over them, thus performing Rao-Blackwellization [2]. This is necessary since the measurement likelihood $p(s_n, o_{n-1}|L^n, z^{n-1})$ cannot be evaluated without knowledge of the landmark locations. Upon performing the marginalization, we obtain (using the notation of (3))

$$p(s_n, o_{n-1}|L^{n,(i)}, z^{n-1}) = \int_{X^n} p(s_n, o_{n-1}|L^{n,(i)}, X^n) \, p(X^n|L^{n,(i)}, z^{n-1}) \tag{8}$$

where X^n is the vector of landmark locations of length n and we have used the chain rule in the integrand. Note that the prior on landmark locations $p(X^n|L^{n,(i)}, z^{n-1})$ can be further factorized into a predictive prior on the location of the current (nth) landmark and the posterior on locations from the previous step

$$p(X^n|L^{n,(i)}, z^{n-1}) = p(X_n|L^{n,(i)}, X^{n-1}, z^{n-1}) \, p(X^{n-1}|L^{n-1,(i)}, z^{n-1}) \tag{9}$$

where X_n is the location of the nth landmark and $p(X^{n-1}|L^{n-1,(i)}, z^{n-1})$ is the posterior on landmark locations from the previous step. Furthermore, the integrand of (8) is equal to the posterior on X^n upto a normalization constant. Hence, we also have a recursive formulation for the posterior on landmark locations $p(X^n|L^{n,(i)}, s^n, o^{n-1})$.

Since storing the posterior on landmark locations at each step aides in the computation of the measurement likelihood, we add this information to each of the particles. However, this is a large continuous space and joint sampling of this space with the space of topologies is not possible. Instead, the posterior is stored in an analytical form, a Gaussian distribution in our case, and updated at each step.

A Rao-Blackwellized Particle Filter (RBPF) maintains the joint posterior over two disparate spaces in exactly the manner described above. Hence, our motivation for using an RBPF in our algorithm is clear. The RBPF maintains the posterior $p(L^n, X^n|s^n, o^{n-1})$ on the joint space of landmark locations and topologies but in a hybrid discrete-continuous form. This posterior can be denoted by a set of samples, each of which also includes an analytical marginal posterior on the landmark locations conditioned on the sample value

$$R_n = \left\{ L^{n,(i)}, w_n^{(i)}, \alpha_n^{(i)}(X^n) \right\}_{i=1}^{N} \tag{10}$$

where $\alpha_n^{(i)}(X^n)$ is the analytic distribution on landmark loca-

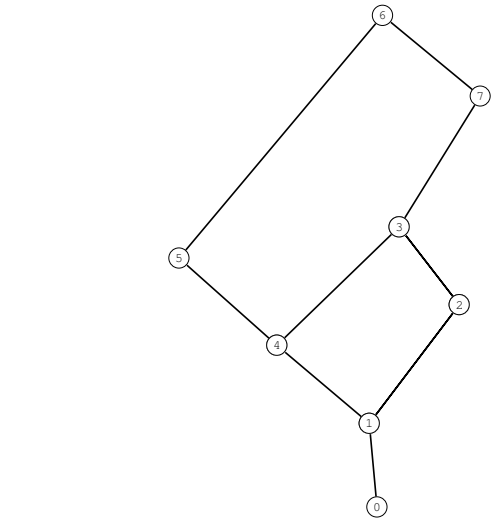

(a)

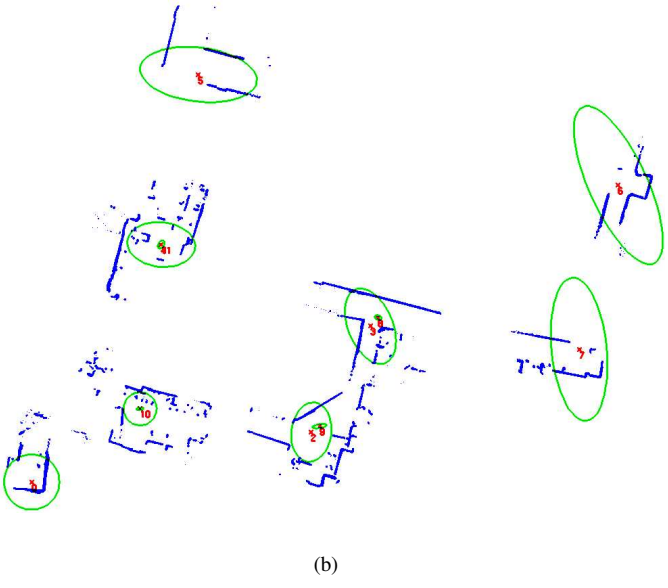

(b)

Fig. 4. A sample from the RBPF that contains (a) a topology and (b) an analytical distribution on the landmark locations in the form of a Gaussian. The red points in (b) are the mean landmark locations while the green ellipses denote marginal covariances.

tions associated with the label sequence $L^{n,(i)}$ defined as

$$\alpha_n^{(i)}(X^n) \triangleq p(X^n|L^{n,(i)}, s^n, o^{n-1}) \quad (11)$$

This analytical distribution can be calculated using Bayes law

$$p(X^n|L^{n,(j)}, z^{n-1}, s_n, o_{n-1}) \propto p(s_n, o_{n-1}|L^{n,(j)}, X^n) \; p(X^n|L^{n,(j)}, z^{n-1}) \quad (12)$$

Note that the second term on the right side of (9) corresponds to $\alpha_{n-1}^{(i)}(X^{n-1})$.

An example of a joint sample from the RBPF is shown in Figure 4.

C. Computing the Importance weights

Computing the importance weights involves calculating the integral in the marginalization step (8). First, consider the measurement likelihood given the landmark locations $p(s_n, o_{n-1}|L^{n,(i)}, X^n, z^{n-1})$. Assuming the independence of the scan and odometry measurements given the landmark measurements, we obtain

$$p(s_n, o_{n-1}|L^{n,(i)}, X^n, z^{n-1}) = p(s_n|L^{n,(i)}, X^n, z^{n-1}) \; p(o_{n-1}|L^{n,(i)}, X^n, z^{n-1}) \quad (13)$$

The scan likelihood $p(s_n|L^{n,(i)}, X^n, z^{n-1})$ is obtained by performing scan matching between the map patches from the landmark locations. The map patches are obtained, in turn, by simply concatenating laser scans from a local area around the landmark as the robot moves through it. To perform the scan matching, we use the scheme of Chen and Medioni [3] that involves point-to-plane matching. The odometry likelihood $p(o_{n-1}|L^{n,(i)}, X^n, z^{n-1})$ is evaluated simply through the use of an odometry model.

The prior on the landmark location $p(X_n|L^{n,(i)}, X^{n-1}, z^{n-1})$ encodes the notion that distinct landmarks do not usually occur close together in the environment. We use the same prior on landmark locations as given in [21]. Topologies which place distinct landmarks close together in location are penalized by this prior. More details can be found in [21].

The weight computation, which is done via the integration in (8), can now be performed in closed form by linearizing around the most likely landmark location and integrating the Gaussian resulting from this. The odometry model is assumed to be Gaussian and the result of the scan matching operation is also a Gaussian distribution. However, the landmark location prior is non-Gaussian. Hence, it is necessary to perform an optimization to find the most likely landmark location and subsequently, linearize around this optimum.

We now have all the components to perform the inference using the RBPF. A summary of the algorithm is provided in Algorithm 1.

Algorithm 1 The RBPF algorithm for inferring PTMs

1) Randomly select a particle $L^{n-1,(i)}$ from the previous time step according to the weights $w_{n-1}^{(i)}$.
2) Propose a new topology sample using the proposal distribution $p\left(L_n^{(j)}|L^{n-1,(i)}\right)$ in (5)
3) Calculate the posterior density on landmark locations $\alpha_n^{(j)}(X^n)$ using Bayes law as in (12) .
4) Calculate the importance weights $w_n^{(j)}$ as the integral over the unnormalized $\alpha_n^{(j)}(X^n)$.

IV. Data Driven Proposal

We suggest the use of a data-driven proposal to overcome a common mode of failure in particle filters, namely the lack of diversity in the particle set as time progresses [5]. The reason for this failure is that many samples fall into regions of low probability and die out during the filtering process. This results not only in the failure to converge to the correct posterior but also in wasted computation, since the algorithm is evaluating the weights of samples that will be lost in any case.

A data-driven proposal overcomes this problem by proposing more samples from regions of high probability so that samples and computation are not wasted. Note that the proposal distribution in (4) does not make use of the current measurement. We rectify the situation in this section by presenting a proposal distribution that uses the odometry to provide more likely samples.

The key idea behind the data-driven proposal is that the odometry likelihood can be incorporated into the proposal distribution while only the scan likelihood is used to compute the importance weights. The measurement likelihood in (1) is thus split into two parts. This split also entails a two-step process for updating the analytic posterior on landmark locations since this posterior needs to be updated using both the odometry and scan measurements.

Starting with the posterior on topologies, we obtain using Bayes Law the likelihood and prior

$$\begin{aligned} p(L^n|s^n, o^{n-1}) &\propto p(L^n|z^{n-1})p(s_n, o_{n-1}|L^n, z^{n-1}) \\ &= p(L^n|z^{n-1})p(o_{n-1}|L^n, z^{n-1}) \\ &\quad p(s_n|L^n, z^{n-1}, o_{n-1}) \end{aligned} \quad (14)$$

where the likelihood is factored into two terms using the chain rule. The prior term can in turn be written using Bayes law as the product of the odometry likelihood and a prior on the current label.

$$p(L^n|z^{n-1}, o_{n-1}) \propto p(L_n|L^{n-1}, z^{n-1})p(L^{n-1}|z^{n-1}) \; p(o_{n-1}|L^n, z^{n-1}) \quad (15)$$

The proposal distribution is taken to be the right hand side of (15), which can be written using the sample representation of (10) as

$$L_n^{(i)} \sim p\left(L_n|L^{n-1,(i)}, z^{n-1}\right) p(o_{n-1}|L^{n,(i)}, z^{n-1}) \quad (16)$$

The form of the predictive label distribution $p(L^n|z^{n-1})$ is the Dirichlet process prior as before. However, the odometry likelihood in (15) is evaluated by marginalization over the

landmark locations

$$p(o_{n-1}|L^n, z^{n-1}) = \int_{X^n} p(o_{n-1}|L^{n,(i)}, X^n, z^{n-1}) \, p(X^n|L^{n,(i)}, z^{n-1}) \tag{17}$$

where the same landmark prior and odometry model are used as in Section III-C. Note that the prior in (17) can be evaluated using the posterior on the landmark locations from the previous by use of the chain rule as in (9).

One drawback of this proposal distribution is the need to perform m optimizations to compute it. These optimizations are required since the integral in (17), evaluated by linearizing around the optimum, needs to be computed for all the possible label values for L_n (except for the case when L_n is a new landmark), which are m in total. However, performing these extra optimizations once per filtering step is still preferable to evaluating the importance weight for all the particles that do not survive when a vanilla proposal is used.

A. Calculating Importance weights

From the target (14) and proposal (15) distributions and the definition of the importance weights (6), we get the expression for the importance weights in this case as

$$w_n^{(i)} \propto p(s_n|L^{n,(i)}, z^{n-1}, o_{n-1})$$

This is evaluated by marginalization over landmark locations

$$w_n^{(i)} \propto \int_{X^n} p(s_n|L^{n,(i)}, X^n, z^{n-1}, o_{n-1}) \, p(X^n|L^{n,(i)}, z^{n-1}, o_{n-1})$$

where the scan likelihood is evaluated using scan matching exactly as in Section III-C, since it is independent of the odometry given the landmark locations. The location prior $p(X^n|L^{n,(i)}, z^{n-1}, o_{n-1})$ is the same as the integrand of (17) upto a normalizing constant and the Gaussian approximation found therein is used again here.

V. Experiments

We validated our algorithms through robot experiments. We used the same datasets in our experiments as used in our previous work [21][22]. The particle filtering was performed using 50 samples and the data-driven proposal was used in all the experiments. A value of 3.0 was used for the Dirichlet prior parameter c. The landmark location prior was used with a value of 10 meters for the penalty radius and 15 for the maximum penalty value. For a description of these parameters and their effect on the inferred posterior, see [21].

The first experiment was conducted using data from an ATRV-mini robot in an indoor setting. A map of the experiment area along with the robot path, which is approximately 100 meters long and passes through twelve landmark locations, is shown in Figure 5. The odometry from the run with the laser scans also plotted is given in Figure 6. The map patches obtained by concatenating scans around the landmark locations are shown in Figure 1.

The result of the filtering using the RBPF algorithm is a joint distribution on topologies and landmark locations.

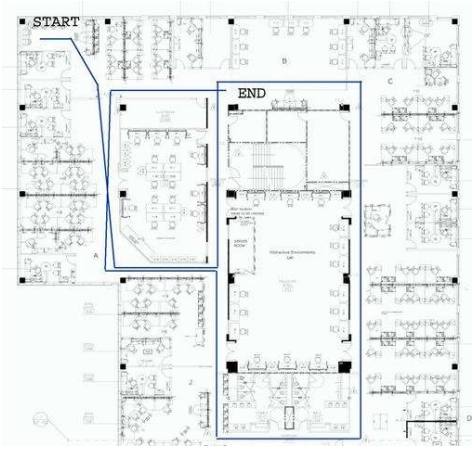

Fig. 5. Schematic of robot path overlaid on a floorplan of the environment for the first experiment.

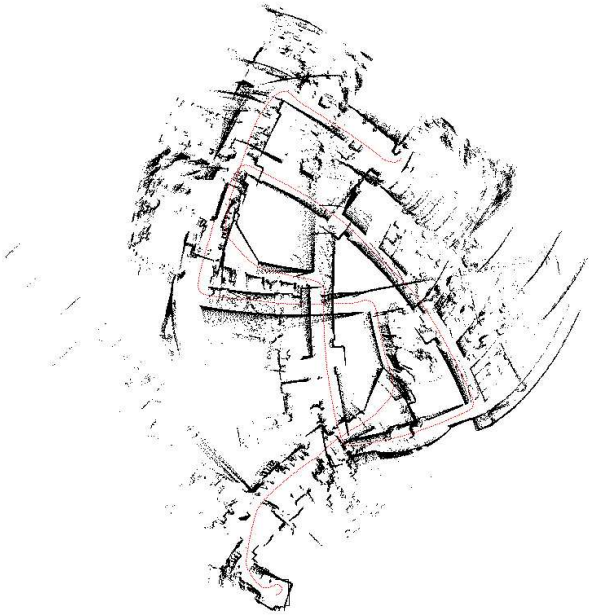

Fig. 6. Robot odometry used in first experiment.

The maximum likelihood sample is shown in Figure 4. The distribution on the landmark locations is displayed in the figure through the marginal covariance ellipses along with the local map patches aligned using scan matching. The corresponding topology, shown in Figure 4(a), is also the ground truth topology and obtains 94% of the probability mass in the posterior. The topology constraints and the inferred landmark locations can be used to produce a global metric map using the global optimization technique of Lu and Milios [14]. The resultant metric map is given in Figure 7. It can be seen that this simple post-processing step produces a globally consistent metric map.

A second experiment was performed in a larger environment (about 60 meters across) to confirm our findings. A floorplan of the test area is shown in Figure 8. The RBPF algorithm computes the PTM that gives the ground-truth topology in Figure 10, 82% of the probability mass. The probability

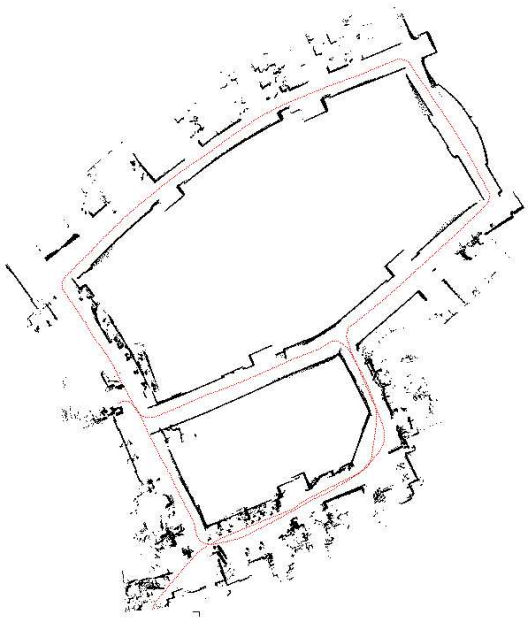

Fig. 7. Global metric map obtained using topological constraints and landmark locations for the first experiment. The robot path is in red.

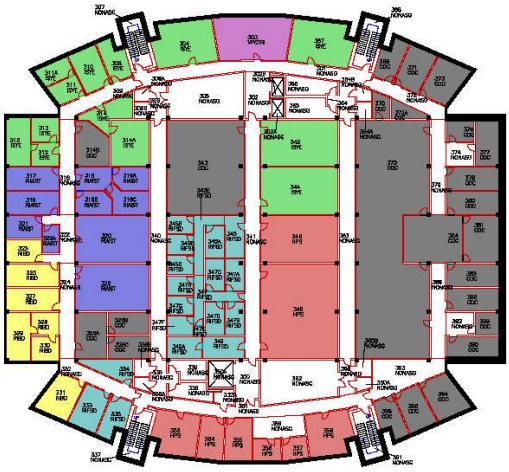

Fig. 8. Floorplan of experimental area for second experiment.

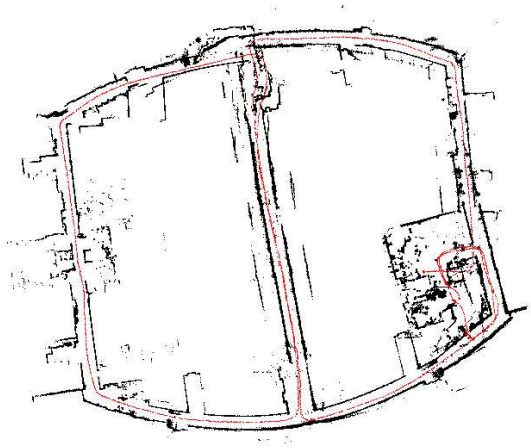

Fig. 9. Metric map obtained using topological constraints for second experiment. The robot path is in red.

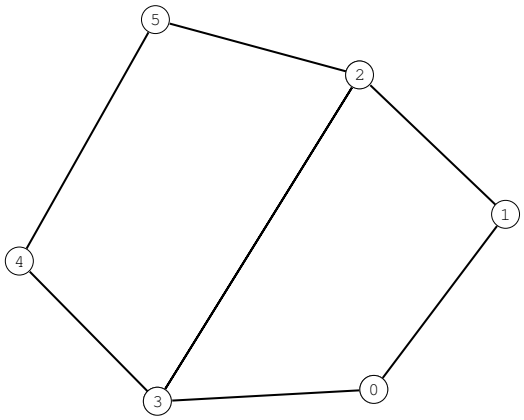

Fig. 10. Ground truth topology for second experiment. This receives 89% of the probability mass in the PTM.

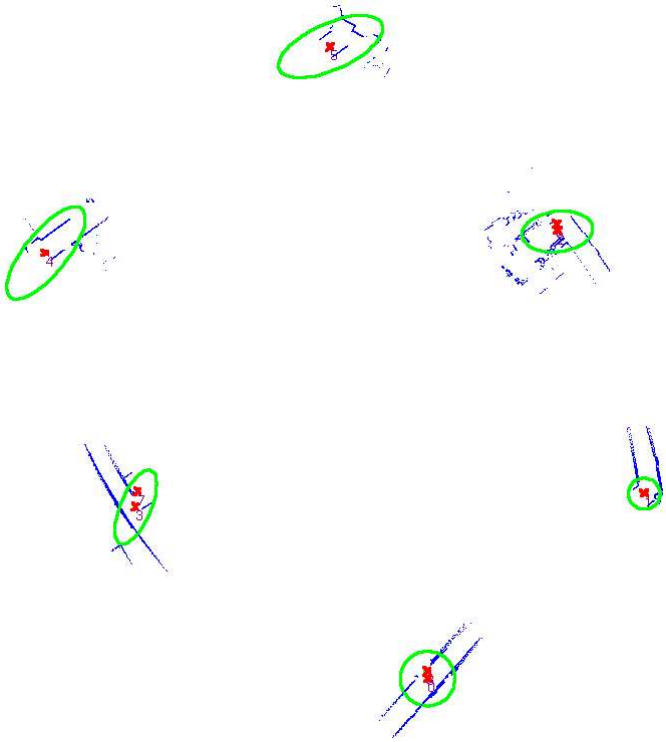

Fig. 11. Maximum likelihood sample from the RBPF for second experiment. The red points are the mean landmark locations while the green ellipses denote marginal covariances.

mass on the ground truth is lower in this case since there is perceptual aliasing around the corners of the building that scan matching is unable to resolve completely. The maximum-likelihood sample with the distribution on landmark locations is shown in Figure 11. The metric map obtained from the Lu-Milios step is given in Figure 9.

VI. Discussion

We presented a novel technique for constructing topological maps in an incremental but robust manner. Our algorithm uses laser range data and odometry in a Rao-Blackwellized Particle Filter (RBPF) setting to perform inference in the joint space of

topologies and landmark locations. The resulting topological maps not only solve the problem of perceptual aliasing but also provide the basis for the construction of global metric maps. The use of a Dirichlet process prior on the landmark labels is also a significant and novel contribution of this work. The use of a data-driven proposal distribution to overcome degeneracy in the particle filter is another contribution. We presented experiments on data gathered from robot runs to validate our algorithm.

An interesting aspect of this work is that computation only needs to be performed at landmark locations. Apart from gathering odometry and detecting for landmarks, the robot is free to perform other tasks between landmarks. Also, since the discrete proposal space at each step is only as large as the number of landmarks observed so far, the particle filter does not require a large number of samples. This fact also contributes to the efficiency of the algorithm.

Recent work by Modayil et. al. [16] is similar to ours in the sense that they too generate an ensemble of topological maps and use them to construct a global metric map. However, they do not provide a probabilistic ordering to their ensemble of maps as the posterior on topologies constructed by our algorithm does. Moreover, since the topological maps generated in that work do not contain sufficient metric information, the process of constructing a global metric map using the topology is a complicated process. By maintaining a joint posterior on the landmark locations in the RBPF, we overcome this drawback.

Other recent work [11][28] has focussed on the creation of hybrid metric-topological maps using local reference frames around landmarks that are connected using the global topological map . By maintaining metric maps only in local frames, these approaches sidestep the problems in creating large scale metric maps. While we do not maintain detailed local maps around landmarks in this work, our approach can easily be extended to generalize such approaches. As an example, each landmark could be associated with a grid map that is updated analytically in the RBPF using a technique similar to [10]. It is future work to implement such a technique that brings the complete power of probabilistic inference to bear on the problem of hybrid metric-topological map construction. It is also future work to integrate the PTM construction with an automatic landmark detector to obtain a complete mapping system on a robot.

ACKNOWLEDGEMENT

This work was funded in part by DARPA under the IPTO/LAGR program (contract #FA8650-04-C-7131).

REFERENCES

[1] D. Blackwell and J.B. MacQueen. Ferguson distributions via polya urn schemes. *Annals of Statistics*, 1:353–355, 1973.

[2] G. Casella and C.P. Robert. Rao-Blackwellisation of sampling schemes. *Biometrika*, 83(1):81–94, March 1996.

[3] Y. Chen and G. Medioni. Object modelling by registration of multiple range images. In *IEEE Intl. Conf. on Robotics and Automation (ICRA)*, pages 2724–2729, 1991.

[4] H. Choset and K. Nagatani. Topological simultaneous localization and mapping (SLAM): toward exact localization without explicit localization. *IEEE Trans. Robot. Automat.*, 17(2):125 – 137, April 2001.

[5] A. Doucet, S.Godsill, and C. Andrieu. On sequential monte carlo sampling methods for Bayesian filtering. *Statistics and Computing*, 10(3):197–208, 2000.

[6] G. Dudek and D. Jugessur. Robust place recognition using local appearance based methods. In *IEEE Intl. Conf. on Robotics and Automation (ICRA)*, pages 1030–1035, 2000.

[7] T. S. Ferguson. A Bayesian analysis of some nonparametric problems. *Annals of Statistics*, 1:209–230, 1973.

[8] R. Gutierrez-Osuna and R. C. Luo. Lola: Probabilistic navigation for topological maps. *AI Magazine*, 17(1):55–62, 1996.

[9] J.-S. Gutmann and K. Konolige. Incremental mapping of large cyclic environments. In *Proc. of the IEEE Intl. Symp. on Computational Intelligence in Robotics and Automation (CIRA)*, pages 318–325, November 2000.

[10] D. Haehnel, W. Burgard, D. Fox, and S. Thrun. A highly efficient FastSLAM algorithm for generating cyclic maps of large-scale environments from raw laser range measurements. In *IEEE/RSJ Intl. Conf. on Intelligent Robots and Systems (IROS)*, 2003.

[11] K. Kouzoubov and D. Austin. Hybrid topological/metric approach to SLAM. In *IEEE Intl. Conf. on Robotics and Automation (ICRA)*, pages 872–877, 2004.

[12] Benjamin Kuipers and Patrick Beeson. Bootstrap learning for place recognition. In *AAAI Nat. Conf. on Artificial Intelligence*, pages 174–180, 2002.

[13] B.J. Kuipers and Y.-T. Byun. A robot exploration and mapping strategy based on a semantic hierarchy of spatial representations. *Journal of Robotics and Autonomous Systems*, 8:47–63, 1991.

[14] F. Lu and E. Milios. Globally consistent range scan alignment for environment mapping. *Autonomous Robots*, pages 333–349, April 1997.

[15] Stephen Marsland, Ulrich Nehmzow, and Tom Duckett. Learning to select distinctive landmarks for mobile robot navigation. *Robotics and Autonomous Systems*, 37(4):241–260, 2001.

[16] J. Modayil, P. Beeson, and B. Kuipers. Using the topological skeleton for scalable global metrical map-building. In *IEEE/RSJ Intl. Conf. on Intelligent Robots and Systems (IROS)*, 2004.

[17] M. Montemerlo, S. Thrun, D. Koller, and B. Wegbreit. FastSLAM: A factored solution to the simultaneous localization and mapping problem. In *AAAI Nat. Conf. on Artificial Intelligence*, 2002.

[18] H.P. Moravec. Sensor fusion in certainty grids for mobile robots. *AI Magazine*, 9:61–74, 1988.

[19] K. Murphy and S. Russell. Rao-Blackwellised particle filtering for dynamic Bayesian networks. In A. Doucet, N. de Freitas, and N. Gordon, editors, *Sequential Monte Carlo Methods in Practice*. Springer-Verlag, New York, January 2001.

[20] A. Nijenhuis and H. Wilf. *Combinatorial Algorithms*. Academic Press, 2 edition, 1978.

[21] A. Ranganathan and F. Dellaert. Inference in the space of topological maps: An MCMC-based approach. In *IEEE/RSJ Intl. Conf. on Intelligent Robots and Systems (IROS)*, 2004.

[22] A. Ranganathan and F. Dellaert. Data driven MCMC for appearance-based topological mapping. In *Robotics: Science and Systems I*, 2005.

[23] E. Remolina and B. Kuipers. Towards a general theory of topological maps. *Artificial Intelligence*, 152(1):47–104, 2004.

[24] F. Savelli and B. Kuipers. Loop-closing and planarity in topological map-building. In *IEEE/RSJ Intl. Conf. on Intelligent Robots and Systems (IROS)*, 2004.

[25] H. Shatkay and L. Kaelbling. Learning topological maps with weak local odometric information. In *Proceedings of IJCAI-97*, pages 920–929, 1997.

[26] K. Sugihara. Some location problems for robot navigation using a single camera. *CVGIP:Image Understanding*, 42(1):112–129, April 1988.

[27] S. Thrun, S. Gutmann, D. Fox, W. Burgard, and B. Kuipers. Integrating topological and metric maps for mobile robot navigation: A statistical approach. In *AAAI*, pages 989–995, 1998.

[28] N. Tomatis, I. Nourbakhsh, and R. Siegwart. Simultaneous localization and map-building: a global topological model with local metric maps. In *IEEE/RSJ Intl. Conf. on Intelligent Robots and Systems (IROS)*, 2001.

[29] N. Tomatis, I. Nourbakhsh, and R. Siegwart. Hybrid simultaneous localization and map building: Closing the loop with multi-hypotheses tracking. In *IEEE Intl. Conf. on Robotics and Automation (ICRA)*, pages 2749–2754, 2002.