

Topological SLAM using Neighbourhood Information of Places

Felix Werner, Frederic Maire, Joaquin Sitte, Howie Choset, Stephen Tully, and George Kantor

Abstract—Perceptual aliasing makes topological navigation a difficult task. In this paper we present a general approach for topological SLAM (simultaneous localisation and mapping) which does not require motion or odometry information but only a sequence of noisy measurements from visited places. We propose a particle filtering technique for topological SLAM which relies on a method for disambiguating places which appear indistinguishable using neighbourhood information extracted from the sequence of observations. The algorithm aims to induce a small topological map which is consistent with the observations and simultaneously estimate the location of the robot.

The proposed approach is evaluated using a data set of sonar measurements from an indoor environment which contains several similar places. It is demonstrated that our approach is capable of dealing with severe ambiguities and, and that it infers a small map in terms of vertices which is consistent with the sequence of observations.

I. INTRODUCTION

A goal of intelligent robotics research is to develop mobile robots which are capable of autonomous navigation. The practicality of such a robot is a function of its ability to use *map-based* navigation to accomplish its mission [1]–[3]. If the environment is unknown, *map-building navigation* is required where the robot performs simultaneous localisation and mapping (SLAM) [2]–[6].

There are two approaches to computing an internal representation of a robot’s environment: *Metric* and *topological* [2], [3]. Metric maps usually capture the geometric properties of the environment [2], [3]. In contrast, a topological map is an abstract and compact representation of the environment that captures *key places* and their *connectivity* for localisation and navigation. Sensory data is used to characterise a place through a *fingerprint*.

For both map inference and localisation, probabilistic approaches have been successfully applied for dealing with the perceptual aliasing problem which lets different parts in the environment appear indistinguishable to the robot. This phenomenon occurs as sensors may supply insufficient data to identify the current state of the world because of measurement uncertainties inherent to robot perception, limited field of view (aperture problem) and repeated structures in the environment. Perceptual aliasing makes it difficult for a

robot to decide when it is visiting a new place or revisiting a memorised place (*loop closing*) [2], [6].

A. Contribution

In this paper we present an approach for topological SLAM using purely allothetic information. We suppose the robot cannot sense any metric information such as odometry or other inertial sensory measurements and we do not make use of the motion actions the robot performed. Our approach presents a general method for topological SLAM and is not restricted to the assumption of purely visual information. However, in order to demonstrate the capability of our approach for dealing with severe perceptual aliasing we use allothetic observations only [7].

In particular, we are interested in the problem of loop closing in environments which contain physically different places which appear to be the same. We approach this problem using Bayesian inference to estimate the posterior distribution on topological maps while simultaneously determining the place the robot currently occupies. The inference method deals with perceptual aliasing by distinguishing similar places on the basis of neighbouring information [8]. Local neighbourhood structures are obtained from the sequence of visited places which is recorded while exploring the environment.

A particular problem is that the number of places is not known in advance. Even if we know the number of places in advance, disambiguation is still difficult unless every place appears different. We propose to solve this ambiguity according to Occam’s razor principle¹ by constructing a small map in terms of places that best explains the sequence of visited places.

B. Related Work

SLAM approaches usually use statistical methods due to the inherent uncertainty and noise in robot perception. Popular are Extended-Kalman-Filter (EKF) based SLAM algorithms [6], [9]. In an EKF the motion model and observation noise are assumed to be independent: the sensory noise is a function of the sensor physics, and is independent of the robot’s motion noise. Monte-Carlo or approaches like FastSLAM [6] are computationally efficient, mapping up to thousands of landmarks while using the EKF for landmark location estimation. These techniques have mainly been implemented in conjunction with metric maps as they require a motion model of the robot to exhibit their functionality.

Felix Werner, Frederic Maire and Joaquin Sitte are with the Faculty of Science and Technology, Queensland University of Technology, Brisbane, QLD 4000, Australia and the NICTA Queensland Lab, St Lucia, QLD 4072, Australia. {felix.werner, frederic.maire, joaquin.sitte}@nicta.com.au

Howie Choset, Stephen Tully and George Kantor are with the Robotics Institute, Carnegie Mellon University, Pittsburgh, PA 15213, USA. {choset@cs, stully@ece, kantor@ri}.cmu.edu

¹“Entia non sunt multiplicanda praeter necessitatem” or “Entities should not be multiplied unnecessarily”. William Occam (1285-1349).

In topological approaches a motion model is difficult to apply because of the abstract representation of the environment. Hence, in topological SLAM approaches the localisation part refers to the *correspondence problem*, that is to recognise previously visited places. Hence, research in topological mapping has for the most part been concerned with a particular aspect of the perceptual aliasing problem. Noisy data sampled at low frequency can cause the robot to perceive fingerprints from distinct places as being non-distinct. In order to simplify the correspondence problem, topological mapping approaches aim to improve the distinctiveness of fingerprints of places [10]–[12]. These approaches do not properly address situations where places are indistinguishable even with perfect sensing.

The robot’s perceptual abilities can be supported using metric information gained from odometry measurements [4], [13]–[16]. However, odometry information is known to be liable to cumulative errors, especially on non-solid and slippery surfaces such as gravel or uneven terrains.

Ranganathan and Dellaert introduced the concept of probabilistic topological maps which, which is a sample-based representation that approximates the posterior distribution over topologies given the available sensor measurements. Mapping is performed through the use of Markov-Chain Monte Carlo based Bayesian inference over the space of all possible topologies [15]. While probabilistic topological maps are a general concept for mapping they are not capable of dealing with repeating structures in the environment.

Werner et al. use neighbourhood information to disambiguate places which appear identical to the robot [8]. Their work examines the approach of using neighbours for disambiguation in terms of applicability and scalability on randomly created, artificial graphs by simulating deterministic observations. In this paper, we exploit the idea of using neighbourhood information to disambiguate places for topological SLAM from a sequence of noisy observations of visited places.

The remainder of this paper is organised as follows: Section II introduces terms and notations which are required to describe our method for topological SLAM. In Section III our algorithm for inferring a topological map and estimating the robot’s location using Bayesian inference is described in detail. Section IV presents results from experiments and the paper is concluded with a discussion section.

II. NEIGHBOURHOOD INFORMATION FOR TOPOLOGICAL MAP INDUCTION

A topological map is represented by a *labelled graph* $\mathbf{G} = (\mathbf{V}, \mathbf{E}, \mathbf{L})$, where *vertices* represent places and *edges* reflect the connectivity between places. The *labels* of vertices refer to fingerprints which characterise the place in terms of sensor data. Each vertex is mapped to a label in \mathbf{L}

$$L : \mathbf{V} \rightarrow \mathbf{L}. \quad (1)$$

If the environment contains perceptual aliasing, several vertices are mapped to the same label so $|\mathbf{L}| < |\mathbf{V}|$, where $|\mathbf{X}|$ denotes the cardinality of set \mathbf{X} .

Definition 1: Vertices x and y of a graph are called *aliases* iff they map to the same label. That is, if $L(x) = L(y)$.

We denote the graph which represents the topology of the environment \mathbf{G}_{env} (*environment graph*) and the robot’s corresponding internal representation \mathbf{G}_{map} (*map graph*). The environment graph is unknown and the only available information about it is a finite *history* $\mathcal{H} = (\mathbf{l}_1, \mathbf{l}_2, \dots) \in \mathbf{L}_{env}^*$ of labels of visited vertices obtained from the traversal of the environment graph (Here, $*$ is the Kleene star).

A. Local Adjacency Information: n -Grams

Our method exploits the neighbourhoods of vertices to disambiguate aliases.

Definition 2: The k -neighbourhood of a vertex $v \in \mathbf{V}$ in a labelled graph $\mathbf{G} = (\mathbf{V}, \mathbf{E}, \mathbf{L})$ is the sub-graph of \mathbf{G} induced by the vertices at distance at most k from v in \mathbf{G} . The parameter k is called the *depth* of the neighbourhood. The vertex set and the connectivity of \mathbf{G}_{env} are unknown. However, the history \mathcal{H} provides information indirectly about \mathbf{G}_{env} . Local neighbourhood information is contained in the history and is accessed through sequences of length n .

Definition 3: A sequence of labels of length n is called *n -gram*.

Consecutively visited vertices correspond to consecutive labels in the history and, reciprocally, consecutive labels in the history correspond to adjacent vertices in \mathbf{G}_{env} . The set of n -grams obtained from a history can be considered as a *feature space* on the history.

Definition 4: The set of all n -grams which can be extracted from the history \mathcal{H} is denoted $Grams_{\mathcal{H}}(\mathcal{H}, n)$.

The number of distinct n -grams which can be extracted from a history of length m is at most $m - n + 1$. Thus, the set of n -grams which is used as input data for topological inference grows linearly with the length of the history. The length of the history can be arbitrary and we assume that it covers all possible n -grams which can be observed from traversing an environment graph.

Given a topological map graph \mathbf{G}_{map} we can generate a history \mathcal{H}_{map} by traversing the map graph. The set of n -grams which can be extracted from all such histories corresponds to a features space on the map graph.

Definition 5: The set of all n -grams which can be obtained by traversing a topological map graph \mathbf{G} is denoted $Grams_{\mathbf{G}}(\mathbf{G}, n)$.

history: $\mathcal{H} = \langle A, B, C, A, E, D, A, B, E, A, C, B, E, D, A, B, C \rangle$
$Grams_{\mathcal{H}}(\mathcal{H}, 3) =$
$\{ \langle A, B, A \rangle, \langle A, B, C \rangle, \langle E, B, A \rangle, \langle A, D, A \rangle, \langle E, D, A \rangle, \langle B, A, B \rangle, \langle D, A, B \rangle, \langle B, C, B \rangle, \langle A, C, B \rangle, \langle B, E, B \rangle, \langle D, E, B \rangle, \langle B, E, A \rangle, \langle C, B, C \rangle, \langle C, B, E \rangle, \langle C, A, C \rangle, \langle C, A, E \rangle, \langle D, A, D \rangle, \langle D, E, D \rangle, \langle A, E, D \rangle, \langle E, B, E \rangle, \langle E, D, E \rangle, \langle E, A, E \rangle, \langle A, C, A \rangle, \langle A, E, A \rangle \}$

TABLE I

AN EXAMPLE OF A POSSIBLE HISTORY h THAT COULD BE OBTAINED FROM THE ENVIRONMENT GRAPH IN FIGURE 1(B) AND THE SET OF 3-GRAMS EXTRACTED FROM THE HISTORY. NOTE, WE ALLOW THE ROBOT TO PERFORM U-TURNS.




Fig. 1. (a) Fictional example environment. The labels A,B,C,D and E denote the view the robot perceives in a particular room. (b) Topological map of the example environment. The vertices are labelled with characters which denote the label (sensor view) of the vertex. The number annotated to each vertex bottom right denotes the index of the vertex. Note the two aliases, labelled A (vertices 1 and 4).

An example environment graph is shown in Figure 1(b). A possible history obtained from traversing this environment graph and, the extracted 3-grams from this history are shown in Table I.

B. n -Consistency

In order to achieve reliable navigation, a robot requires an internal representation that exhibits the properties of the environment with respect to the selected representation (e.g. topological). Consequently, it is desirable that a map graph is isomorphic with the environment graph.

A necessary condition for the map graph to be isomorphic to the environment graph is that there is a bijective mapping such that each k -neighbourhood of the map graph corresponds to a k -neighbourhood of the environment graph and each k -neighbourhood of the environment graph corresponds to a k -neighbourhood of the map graph. However, in our case it is not possible to compare k -neighbourhoods of the map graph directly with k -neighbourhoods in environment graph as the latter is unknown. Consequently we propose to measure the n -consistency of graphs in their feature spaces; that means, the sets of n -grams of the graphs.

Definition 6: Two sets of n -grams Γ_0 and Γ_1 are n -consistent, where $n > 1$, iff $\Gamma_0 = \Gamma_1$.

According to Definition 6 two entities which induce n -grams, such as graphs or histories, are n -consistent if they share the same set of n -grams. The n -consistency concept realises the idea of k -neighbourhoods of vertices (see Definition 2) on histories if $n = 2k + 1$.

For noisy data we measure the degree of n -consistency of two sets of n -grams Γ_0 and Γ_1 using the Hausdorff distance

$$d_H(\Gamma_0, \Gamma_1) = \max(\max_{\alpha \in \Gamma_0} \min_{\beta \in \Gamma_1} d(\alpha, \beta), \max_{\beta \in \Gamma_1} \min_{\alpha \in \Gamma_0} d(\beta, \alpha)). \quad (2)$$

The smaller the Hausdorff distance the more n -consistent are Γ_0 and Γ_1 . The distance of two n -grams α and β is computed using the maximum norm

$$d(\alpha, \beta) = \|\alpha - \beta\|_\infty = \max_{k=0, \dots, n-1} (|\alpha_k - \beta_k|). \quad (3)$$

Thus, the distance between two sets of n -grams is determined by the largest distance of two fingerprints which are mapped to the same vertex in a map graph.

C. n -Consistency and Map Size

In parametric methods there may be various candidate models, each with a different number of parameters to represent the data. The likelihood of the model to represent the data is increased when the number of parameters in the model is increased. In our case, a map graph which consists of one vertex for each n -gram in $\text{Grams}\mathcal{H}(\mathcal{H}_{env}, n)$ is consistent according to Definition 6 but is inappropriate for navigation as too many vertices would be required.

III. TOPOLOGICAL SLAM

In this section we describe our method for topological SLAM from a sequence of visited places. We suppose the robot has explored an environment and recorded a history $\mathcal{H}^{1:T} = \mathbf{I}^1, \dots, \mathbf{I}^t$ of fingerprints of visited places.

We suppose the robot has some prior knowledge about the environment given in a set $\Gamma_{env} = \text{Grams}\mathcal{H}(\mathcal{H}_{env}, n)$ of n -grams which may be obtained in an earlier exploration by traversing every intersection from every possible direction. The set Γ_{env} only contains information about the connectivity but does not comprise information about the number of places in the environment.

In the following we describe a method for topological map inference using a particle filter which embeds the disambiguation using neighbourhood information.

A. Map Likelihood Estimation using a Particle Filter

Particle filters are sequential Monte-Carlo methods used for Bayesian model inference. Bayesian filters probabilistically estimate a dynamic system's state from noisy observations. Here, the system state $\mathbf{s}^t = \{\mathbf{G}_{map}^t, p^t\}$ of the world comprises the topological map \mathbf{G}_{map}^t and the place p^t the robot visits. We are interested in computing the posterior PDF (probability density function) $P(\mathbf{s}^t | \mathcal{H}^{1:t}, \Gamma_{env})$ over the state at the time step t given all the measurements up to the current time. Particle filters model the PDF on the state in a Monte-Carlo fashion using a collection $\{\mathbf{s}_i^t, w_i^t\}_{i=1}^N$ of weighted *particles*. Particles are denoted \mathbf{s}_i^t and the w_i^t are non-negative weights, called *importance factors* which are normalised such that

$$\sum_{i=1}^N w_i^t = 1. \quad (4)$$

In order to perform topological SLAM we need to recursively compute the density $P(\mathbf{s}^t | \mathcal{H}^{1:t}, \mathbf{\Gamma}_{env})$ at each time step. This is done in two phases:

1) *Prediction Phase*: In the prediction phase a transition model is commonly used to predict the new state of the system at $t + 1$ using a predictive PDF $P(\tilde{\mathbf{s}}^{t+1} | \mathcal{H}^{1:t}, \mathbf{\Gamma}_{env})$. The transition model is specified as a conditional density $P(\tilde{\mathbf{s}}^{t+1} | \mathbf{s}^t, \mathbf{u}^t)$ and is dependent on the previous state \mathbf{s}^t (Markov) and a known control input \mathbf{u}^t . Because we assume the motion of the robot is not known it is not possible to predict the place the robot visits next. Hence, we model the transition density $P(\tilde{\mathbf{s}}^{t+1} | \mathbf{s}^t, \mathbf{u}^t)$ by generating a collection of predictions $\{\tilde{\mathbf{s}}_{i,k}^{t+1}\}_{k=1}^{K+1}$ for each particle \mathbf{s}_i^t , where k indicates the vertex the new observation is predicted to correspond to and K refers to the number of vertices contained in the map graph particle \mathbf{s}_i^t holds.

An edge is added between the vertex the robot currently visits and the vertex the robot is predicted to visit next in case that particular adjacency does not exist yet.

2) *Update Phase*: In the update phase a measurement model is used to incorporate information from the sensors to obtain the posterior PDF $P(\mathbf{s}^{t+1} | \mathcal{H}^{1:t+1}, \mathbf{\Gamma}_{env})$. We assume the measurement model is given in terms of a measurement likelihood $P(\mathcal{H}^{t+1} | \tilde{\mathbf{s}}_i^{t+1}, \mathbf{\Gamma}_{env})$ which expresses the likelihood that the topological map and the location comprised in the prediction $\tilde{\mathbf{s}}_i^t$ corresponds to the true map and location. The posterior density over \mathbf{s}^{t+1} is then obtained using Bayes theorem:

$$P(\mathbf{s}^{t+1} | \mathcal{H}^{1:t+1}, \mathbf{\Gamma}_{env}) \propto P(\mathcal{H}^{t+1}, \mathbf{\Gamma}_{env} | \tilde{\mathbf{s}}^{t+1}) P(\tilde{\mathbf{s}}^{t+1} | \mathcal{H}^{1:t}, \mathbf{\Gamma}_{env}). \quad (5)$$

In our particle filter approach the measurement likelihood is computed by weighting the samples

$$w_i^t = P(\mathcal{H}^{t+1}, \mathbf{\Gamma}_{env} | \tilde{\mathbf{s}}_i^{t+1}) = P(\mathcal{H}^{t+1} | \tilde{\mathbf{s}}_i^{t+1}) P(\mathbf{\Gamma}_{env} | \tilde{\mathbf{s}}_i^{t+1}) P(\tilde{\mathbf{s}}_i^{t+1}). \quad (6)$$

The factor $P(\mathcal{H}^{t+1} | \tilde{\mathbf{s}}_i^{t+1})$ computes the likelihood of the measured label \mathcal{H}^{t+1} and the label $\tilde{\mathbf{I}}_k$ of the vertex k which refers to the predicted new location $\tilde{\mathbf{p}}_i^{t+1}$ of the robot in particle $\tilde{\mathbf{s}}_i^{t+1}$ to be identical

$$P(\mathcal{H}^{t+1} | \tilde{\mathbf{s}}_i^{t+1}) = \exp \left(- \left(\frac{\mathcal{H}^{t+1} - \tilde{\mathbf{I}}_k}{\sigma_l} \right)^2 \right) \quad (7)$$

where σ_l denotes a weighting factor.

The second factor in Equation 6 expresses the likelihood that the graph in particle $\tilde{\mathbf{s}}_i^{t+1}$ is n -consistent with the information given in the history. Using Equation 2 the consistency likelihood is computed with

$$P(\mathbf{\Gamma}_{env} | \tilde{\mathbf{s}}_i^{t+1}) = \exp \left(- \left(\frac{d_H(\mathbf{\Gamma}_{env}, \tilde{\mathbf{\Gamma}}_i^{t+1})}{\sigma_c} \right)^2 \right) \quad (8)$$

where $\tilde{\mathbf{\Gamma}}_i^{t+1} = Grams_G(\tilde{\mathbf{G}}_i^{t+1}, n)$. In order to address the conflicting interests of inferring the a map in terms of vertices that explains the history well, we seek to find a compromise




Fig. 2. Penalty function to avoid maps containing many vertices.

that the data and map space allow. Hence, the last term in Equation 6 penalises map graphs which contain vertices with similar labels

$$P(\tilde{\mathbf{s}}_i^{t+1}) = \prod_{a=1}^A \prod_{b=1}^A \left(1 - \phi \exp \left(- \left(\frac{\tilde{\mathbf{I}}_a - \tilde{\mathbf{I}}_b}{\sigma_l} \right)^2 \right) \right) \quad (9)$$

where ϕ weights the influence of the penalty and $\tilde{\mathbf{I}}_a, \tilde{\mathbf{I}}_b \in \tilde{\mathbf{L}}_i^{t+1}$, $A = |\tilde{\mathbf{V}}_i^{t+1}|$ and $\tilde{\mathbf{L}}_i^{t+1}$ and $\tilde{\mathbf{V}}_i^{t+1}$ refer to the set of labels and vertices of the predicted map graph comprised in particle $\tilde{\mathbf{s}}_i^{t+1}$, respectively. The penalty function for $\phi = 1$ is displayed in Figure 2.

The posterior distribution on topological maps is computed by drawing N samples from the proposal distribution.

B. Localisation

The place p^t the robot occupies is implicitly estimated whenever the map graph is updated with a new observation. The vertex whose label is updated or additionally introduced using the observation indicates the new location of the robot. If a new vertex is introduced the robot is hence located at the place which corresponds to that vertex. The location, in turn, is used to guide the mapping process by introducing adjacencies between the current and the previous place occupied.

IV. RESULTS FROM EXPERIMENTS

Our experimental set up covers an indoor office environment area of about 20,000 square meters (Wean Hall at Carnegie Mellon University), see Figure 3. The experimental platform uses a panoramic camera and an ultrasonic sensor array to acquire information about the environment.

A. Place Identification using GVG

The experimental platform traverses the environment using the generalised Voronoi graph (GVG) strategy developed by Choset and Nagatani [4]. It is based on the Voronoi diagram which is a special kind of decomposition of a metric space into segments and nodes determined by distances to a specified discrete set of objects in the space. Our robot measures distances using sonar readings.

Segments capture the points in the plane that are equidistant to two sites. Travelling along the Voronoi segments, the robot can keep in the middle of corridors while exploring the environment. The Voronoi nodes are the points equidistant to three (or more) obstacles.

In our system, the Voronoi nodes define the places in the environment and the range measures of the sonar sensors




Fig. 3. The floor plan of Wean Hall Floor 6 at Carnegie Mellon University. Embedded is the topological graph (vertices and their connectivity) that reflects the ground truth of the topological map that we wish the SLAM algorithm to infer. Moreover, the travelled path of the robot is displayed. The circles show the distance measure from the locus point to the closest obstacles which we use as fingerprints of the places.

to the closest obstacles at these locus points represent the fingerprints of the places. The fingerprint may be interpreted as the radius of the largest obstacle free disc that can be drawn around the locus points of Voronoi nodes.

We have conducted several exploration runs and recorded a total of 105 fingerprints of places and stored them in a data base (see Table II). It is apparent that the environment in conjunction with the very basic description of places contains numerous topological ambiguities. In fact, the environment contains eleven places but only four categories of fingerprints, resulting in a challenge for topological mapping and localisation algorithms. For example, places 1 and 2 appear similar to the robot as do places 3 and 4 or places 2, 5, 6, 7 and 10.

B. Results

Given the data base and the ground truth environment graph, we can simulate arbitrary traversals of the environment. The robot starts at an initial vertex and selects an arbitrary adjacent vertex as next place. According to the vertex the robot occupies, a random observation from the data base is sampled. For the following evaluations, 500 paths of length 100 were generated. Each path represents an exhaustive exploration of the environment. For each path the set Γ_{env} of n -grams is derived before starting the algorithm. Note, it is actually not necessary to assume an exhaustive exploration of the environment as the inferred map is a representation of the environment which is consistent to the measurements at a certain time.

It was shown by Werner et al. that environments which contain many repeated places require a higher level of consistency [8]. Hence, for the following evaluations, the mapping part of our approach aims for 5-consistency between the inferred topological map and the information given in the history of observations.

Our approach to topological mapping aims to build an internal representation which is consistent with the information given in the history. The consistency is measured through the Hausdorff distance (Equations 2 and 3) between the sets of n -grams which are derived from the history and the inferred map graph. Figure 4 shows a histogram of the 5-consistency measure of the inferred maps of the simulated random traversals. It can be seen that most of the inferred maps are very consistent with the information from the history. The

maximum consistency error of the fingerprints of vertices contained in the map and the corresponding observations is mostly less than 10cm which can be explained with the inherent uncertainties associated with sensory perception. The rare outliers may occur when the mapping process is misled so that the inferred map is inconsistent with the observations.

As mentioned earlier, a topological map that maximises the consistency between the inferred map and the history would contain a vertex for every observation. We penalise topological maps which contain vertices with similar labels to obtain small map graphs in the number of vertices. Figure 5 shows the number of vertices of the inferred maps using $\phi = 0.1$ as penalty parameter. The clear peak of the histogram shows that most of the maps our algorithm induced require eleven vertices what corresponds to the ground truth (see Figure 3). It seems the algorithm does not infer smaller maps as they would violate the consistency criterion and hence yield low probability. Larger maps indicate an additionally introduced vertex so the map erroneously contains two vertices which represent the same place in the environment. Erroneously introduced vertices occur due to measurement noise and, as the sampling of new map candidates from a posterior distribution can sample map candidates which are not optimal and so create larger maps. At the moment our system cannot recover from this problem. In general, a smaller penalty parameter on the number of vertices results in larger maps whereas a high penalty parameter decreases the map size but may increase the consistency error.

The overall goal in topological mapping is to build an internal representation which is isomorphic to the environment. Here, we investigate whether the inferred map graphs are isomorphic to the environment graph in order to measure the quality of the proposed approach for topological mapping. We found all map graphs with the same number of vertices as the environment graph to be isomorphic to the environment graph. For the evaluations we used 10 samples to represent the posterior distribution on topological maps. Increasing the number of samples results in a higher ratio of isomorphic

id	Recordings	Mean [m]	Var
1	13	1.1713	0.0121
2	11	1.3600	0.0069
3	11	2.1641	0.0179
4	9	2.1624	0.0077
5	9	1.3350	0.0107
6	7	1.2695	0.0255
7	8	1.3104	0.0008
8	7	1.4983	0.0024
9	10	1.6768	0.0134
10	10	1.3537	0.0009
11	10	1.1606	0.0110

TABLE II
DATA BASE OF SONAR READINGS RECORDED FROM SEVERAL TRAVERSALS OF THE EXPERIMENTAL ENVIRONMENT. THE IDS OF THE PLACES CORRESPOND TO FIGURE 3.




Fig. 4. The consistency of the inferred topological maps and the corresponding simulated histories is displayed. It is to see that most of the inferred maps have only small consistency errors which occur due to measurement noise.




Fig. 5. The number of vertices of the induced maps are shown for the simulated histories. A clear peak for 11 vertices is to see what corresponds to the environment graph (see Figure 3).

maps, however, requires more computational resources. The results support the strategy of the proposed algorithm to use the current position estimation with the new observation to map the connectivity of the environment.

Our method for topological SLAM implicitly localises the robot. The particle filter technique we use requires to predict the next observation and hence essentially predicts the location the robot visits next. Reliable localisation is crucial as the location directly governs the connectivity inference. The localisation performance of our method is hence implicitly evaluated through the resulting topological maps. For the resulting maps which are isomorphic to the environment graph the location of the robot must always be estimated correctly.

V. DISCUSSION

In this paper we have proposed a Bayesian approach for topological SLAM that does not require any motion model or metric information, but uses a history of noisy measurements from visited places only. While we restricted ourselves to sonar measurements only, the presented approach for topological SLAM allows to include further information sources which may help to increase the quality of both, induced maps and localisation of the robot.

We use a particle filter based SLAM approach to deal with both measurement noise perceptual aliasing. The posterior distribution on topologies aims to maintain consistency with the observed data while minimising the number of vertices contained in the map. The consistency between a topological map and the observations is measured using the Hausdorff distance.

Experiments in an indoor environment which is subject to severe ambiguities due to repeated structures demonstrate the

merit of the idea to use neighbourhood clues in order to disambiguate otherwise identical vertices. Our approach mostly infers topological maps with only small inconsistencies with respect to the data. Moreover, many of the resulting maps are isomorphic to the environment graph.

While we used sonar data only for evaluation, our approach presents a general frame work for topological map inference and allows to incorporate odometry observations or information about the motion actions of the robot for enhancing the transition model of the robot.

VI. ACKNOWLEDGEMENTS

NICTA is funded by the Australian Government as represented by the Department of Broadband, Communications and the Digital Economy and the Australian Research Council through the ICT Centre of Excellence program.

REFERENCES

- [1] S. Thrun, "Robotic mapping: A survey," in *Exploring Artificial Intelligence in the New Millennium*, G. Lakemeyer and B. Nebel, Eds. San Francisco, CA, USA: Morgan Kaufmann Publishers Inc., July 2002, pp. 1–35.
- [2] D. Filliat and J.-A. Meyer, "Map-based navigation in mobile robots: I. a review of localization strategies," *Cog. Sys. Res.*, vol. 4, no. 4, pp. 243–282, December 2003.
- [3] J.-A. Meyer and D. Filliat, "Map-based navigation in mobile robots: II. a review of map-learning and path-planning strategies," *Cog. Sys. Res.*, vol. 4, no. 4, pp. 283–317, December 2003.
- [4] H. Choset and K. Nagatani, "Topological simultaneous localization and mapping (slam): Toward exact localization without explicit localization," *IEEE Trans. Robot. Automat.*, vol. 17, no. 2, pp. 125–137, 2001.
- [5] J. J. Leonard and H. F. Durrant-Whyte, "Simultaneous map building and localization for an autonomous mobile robot," in *IROS*, vol. 3. Osaka, Japan: IEEE, November 1991, pp. 1442–1447.
- [6] S. Thrun, W. Burgard, and D. Fox, *Probabilistic Robotics*. Cambridge, MA: MIT Press, 2005.
- [7] F. Werner, F. Maire, and J. Sitte, "Topological slam using fast vision techniques," in *FIRA*, ser. LNCS, no. 5744. Incheon City, Korea Republic: Springer Verlag Berlin Heidelberg, August 2009, pp. 187–196, to appear.
- [8] F. Werner, C. Gretton, F. Maire, and J. Sitte, "Induction of topological environment maps from sequences of visited places," in *IROS*. Nice, France: IEEE, September 2008, pp. 2890–2895.
- [9] M. W. M. G. Dissanayake, P. Newman, S. Clark, H. F. Durrant-Whyte, and M. Csorba, "A solution to the simultaneous localization and map building (slam) problem," *IEEE Trans. Robot. Automat.*, vol. 17, no. 3, pp. 229–241, June 2001.
- [10] A. Tapus and R. Siegwart, "Incremental robot mapping with fingerprints of places," in *IROS*. Edmonton, AB, Canada: IEEE, August 2005, pp. 2429–2434.
- [11] C. Valgren, T. Duckett, and A. Lilienthal, "Incremental spectral clustering and its application to topological mapping," in *ICRA*. Roma, Italy: IEEE, April 2007, pp. 4283–4288.
- [12] H. Andreasson and T. Duckett, "Topological localization for mobile robots using omni-directional vision and local features," in *IAV*. Lisbon, Portugal: Elsevier, July 2004, pp. 3348 – 3353.
- [13] P. Lamon, A. Tapus, E. Glauser, N. Tomatis, and R. Siegwart, "Environmental modeling with fingerprint sequences for topological global localization," in *IROS*, vol. 4. Las Vegas, NV, USA: IEEE, October 2003, pp. 3781–3786.
- [14] F. Werner, J. Sitte, and F. Maire, "Visual topological localisation and mapping using colour histograms," in *ICARCV*. Hanoi, Vietnam: IEEE, December 2008.
- [15] A. Ranganathan and F. Dellaert, "Inference in the space of topological maps: An mcmc-based approach," in *IROS*, vol. 2. Sendai, Japan: IEEE, September - October 2004, pp. 1518–1523.
- [16] A. De, J. Lee, and N. J. Cowan, "Toward slam on graphs," in *WAFR*, Guanajuato, México, December 2008.