A New Approach for Large-Scale Localization and Mapping: Hybrid Metric-Topological SLAM

Jose-Luis Blanco, Juan-Antonio Fernández-Madrigal, Javier Gonzalez
Dept. of System Engineering and Automation
University of Malaga
Málaga, Spain
{jlblanco,jafma,jgonzalez}@ctima.uma.es

***Abstract*—Most successful works in Simultaneous Localization and Mapping (SLAM) aim to build a metric map under a probabilistic viewpoint employing Bayesian filtering techniques. This work introduces a new hybrid metric-topological approach, where the aim is to reconstruct the path of the robot in a hybrid continuous-discrete state space which naturally combines metric and topological maps. Our fundamental contributions are: (i) the estimation of the *topological path*, an improvement similar to that of Rao-Blackwellized Particle Filters (RBPF) and FastSLAM in the field of metric map building; and (ii) the application of grounded methods to the abstraction of topology (including loop closure) from raw sensor readings. It is remarkable that our approach could be still represented as a Bayesian inference problem, becoming an extension of purely metric SLAM. Besides providing the formal definitions and the basics for our approach, we also describe a practical implementation aimed to real-time operation. Promising experimental results mapping large environments with multiple nested loops (~30.000 m^2, ~2Km robot path) validate our work.**

Index Terms*—Mobile robots, Large-scale maps, Loop closure, Rao-Blackwellized Particle Filters, SLAM.

I. Introduction

Simultaneous Localization and Mapping (SLAM) has being intensively studied by researchers in the last decade, leading to approaches that can be classified into three well-differentiated paradigms depending on the underlying map structure: metric ([5],[12],[24],[26]), topological ([22],[23]), or hybrid representations ([7], [15],[28]).

In this work we focus on the formulation of hybrid metric-topological SLAM in terms of Bayesian estimation from all the available robot actions and observations. Thus, our major contribution consists of a new formulation of hybrid SLAM as the estimation of the sequence of areas the robot traverses (topological part) and the local pose of the robot within those areas (metric part), providing an estimate of the spatial relation between elements, not their absolute positions [9]. Our proposal is grounded on the fact that the robot will be always at *one* area, in whose metric scope SLAM is a solved problem by standard methods, either Extended Kalman Filters [5] or RBPF [6]. As long as the size of the areas (local maps) is kept bounded, so does the complexity of the local SLAM method. In this paper we discuss issues such detecting when the robot goes out of the current area, enters a new one, or reenters a previously known one (loop closure). In our proposal two estimation processes are carried out concurrently: (i) the robot pose relative to the current area (metric path), and (ii) the sequence of areas the robot goes through (topological path). Under this perspective, loop closure becomes finding a partition in the sequence of all the areas in the map [22].

The meaning of the topological map in this work strongly differs from the one considered in many other works. In the literature we can find works that consider distinctive places as nodes ([4],[15]), while others cut the map into disjoint areas ([7],[28]). Our model is closer to those of appearance-based maps ([2],[30]), where topological nodes are the result of abstracting low-level robot observations gathered at a given area. Actually, the size of areas will be automatically determined by the nature of sensors, more concretely, by the covisibility between observations [2]. As an example, observations gathered by a laser scanner within a room may be grouped into a single area, but using a narrow field-of-view camera may generate a number of areas instead. Furthermore, partitioning observations, as distinct from the physical space, implies that the same location may be assigned different relative coordinates, one for each area in the topological map.

In addition to providing a unified theoretical support for hybrid metric-topological SLAM (which we call HMT-SLAM), in this work we propose a practical implementation framework. This system processes local metric information in real-time by means of a particle filter with a constant-time complexity, while the topological structure of the environment is estimated in an anytime fashion.

We claim that the present work is a promising base for integrating many previous separate contributions, such as:

- Rao-Blackwellized Particle Filters (RBPF) for mapping, which suffer major problems when dealing with large or nested loops, requiring a dynamic number of particles [8] or artifacts to prevent the loss of diversity [25], respectively. These problems are the result of map building under a global-coordinate approach.
- In the context of occupancy grid map building, advanced techniques ([10],[11]) are required to reduce the memory requirements of RBPF mapping, due to the maintenance of a global map in each particle. In our approach each particle keeps a map of the current area only, hence achieving improved scalability with a great reduction in storage requirements.
- Appearance-based map partitioning methods ([2],[30]) have never been integrated before into a hybrid SLAM framework, whereas they naturally fit into the induction of topological areas from metric observations.

- Global localization (the robot "awakening" problem), can be managed in a more efficient manner in our hybrid state-space than within a global metric map. The robot can first localize in which area it is, and then try to estimate its metric within that area. However, this issue will not be addressed here for the sake of brevity.

The rest of this paper is structured as follows. In section II we examine previous works related to both metric and topological SLAM. Next, we provide the probabilistic foundations of our approach, while some of its relevant elements are discussed in section IV. A practical system that implements our ideas is presented in section V, and experimental results with real robots in large-scale scenarios are discussed in section VI. Finally, some conclusions and future work are outlined.

II. Previous Research

In the following we briefly discuss previous works in the fields of metric, topological, and hierarchical mapping, highlighting their relation with the present paper.

Metric approaches ([12],[17],[19],[20],[24],[26]) aim to reconstruct the spatial arrangement of map elements, in the form of landmark maps [19], occupancy grids [20], or sets of range scans [17] (please refer to [27] for a more detailed classification). Although some non-probabilistic methods have been proposed to build metric maps ([12],[17]), the vast majority of works on metric SLAM rely on a probabilistic representation of the robot pose and the map, where Bayesian filtering estimates the corresponding probability distributions [5]. The hardest problem in these methods is data association (that is, establishing correspondences among observations and the map) [21], a problem that aggravates when the robot closes a loop: a previously known place is revisited through a new and unknown path. Establishing wrong correspondences greatly compromises the consistency of estimated maps. Since the uncertainty in the robot pose and the map increases as the robot explores new areas, the hardness of finding the correct data association increases with the scale of maps. Recently, this problem has been successfully addressed from a new viewpoint: estimating the whole robot *path*, instead of only the most recent *pose*. We can then apply a convenient factorization, called Rao-Blackwellization in Estimation Theory [6], that has enabled the mapping of relatively large-sized environments with both occupancy grids (Rao-Blackwellized Particle Filters [10]) and landmark maps (FastSLAM [18]). However, the number of particles required to close a loop increases with its length, which may eventually turns into a storage capacity limitation since each particle must carry a hypothesis of the whole map.

Building a topological map is an attractive alternative to metric maps. Among other properties, they have reduced storage requirements and can be easily integrated with symbolic planning. Although Bayesian estimation has been reported for these maps [22], it is assumed there that the robot can detect whether it is close to one of a set of "distinctive" places, which are represented as nodes in the map. We think this is a too restrictive assumption where the diversity of metric information is lost. A more appealing approach is to consider hybrid maps, where topological nodes contain local metric information ([3],[7],[14],[15],[16]). However, loop closure for these maps in previous works has been considered only under the metric point-of-view, i.e. by finding the global coordinates transformation compatible with the loop closure [7]. For this reason the association problem for observations remains being a major issue, which can be simplified by inferring probabilistic topological loop closure hypotheses [22]. Our approach turns the landmark-to-landmark data association problem into a node-to-node one, which we claim is an easier problem due to its reduced ambiguity. Additionally, to the best of our knowledge no previous work has proposed the probabilistic estimation of the topological path followed by the robot, which may be seen as the dual of Rao-Blackwellization for metric mapping.

III. Probabilistic Foundations Of HMT-SLAM

The problem of metric SLAM is stated as to simultaneously estimate the map m and the robot pose x_t at any given instant of time t. Set out as a Bayesian filtering problem conditioned to the sequence of robot actions $u^t=\{u_1,\ldots,u_t\}$ and observations $z^t=\{z_1,\ldots,z_t\}$ [5], the probability distribution to be estimated is [1]:

$$p\left(x_t, m \middle| u^t, z^t\right) \tag{1}$$

Under this formulation of the problem, all the observations z_t depend on the whole map m. This is in contrast with the "locality" of real observations, which typically catch a small part of the environment at once.

Our proposal for a hybrid SLAM builds upon the assumption that the map can be conveniently divided into a set of n metric sub-maps $\{{}^{i}\mathcal{M}\}_{i=1..n}$, which we will call *areas*. Each area has its own coordinate reference frame, while *edges* $\{{}^{ab}\Delta\}$ between sub-maps define the transformation between different frames leading to the topological view of the map. Thus, a hybrid map is a 2-tuple:

$$m = \left\langle \left\{ {}^{k}\mathcal{M} \right\}_{k=1..n}, \left\{ {}^{ab}\Delta \right\}_{\substack{a=1..n \\ b=1..n}} \right\rangle \tag{2}$$

Accordingly, the robot pose becomes a hybrid metric-topological (HMT) variable, stated as $s_t=(x_t,\gamma_t)$, where the discrete part γ_t identifies a sub-map (area), and the continuous x_t is the pose relative to the coordinate reference of the area ${}^{\gamma_t}\mathcal{M}$.

Given the above definition of a HMT map m, we state the HMT-SLAM problem as the estimation of the following distribution:

$$p\left(s^t, m \middle| u^t, o^t\right) \tag{3}$$

where $o^t=\{o_1,\ldots,o_t\}$ is the sequence of hybrid observations $o_t=(z_t, \psi_t)$. The purpose behind defining o_t as hybrid is to conveniently separate metric observations z_t, and area dependant, qualitative observations ψ_t (this division is especially useful when facing the problem of global

[1] For clarity, throughout this paper we will denote sequences of variables of the form $x_{1:k}=\{x_1,\ldots,x_k\}$ as x^k.

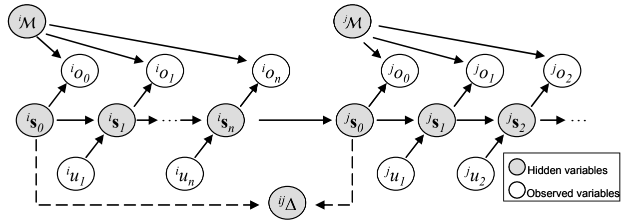

Fig. 1 The graphical model for HMT-SLAM. Here segments of the robot path are conditionally independent given the starting pose from each segment. The relative pose between areas ${}^{ij}\Delta$ is a random variable but defined as an analytical function of robot poses.

localization). The estimation of the topological and metric *paths* enables the Rao-Blackwellization of the joint map-path estimation [6]. It must be highlighted the clear advantage of this model to face the hard problem of loop closure: while metric SLAM methods aim to compute the exact coordinates of the robot after traversing a cycle (i.e. computing continuous variables), closing the loop in the topological scope is equivalent to establishing a partition in the (discrete) space of topological areas. This issue was rigorously discussed elsewhere [22]. One could argue that the metric elements in the HMT map (edges ${}^{ab}\Delta$ between areas) still have to be considered. However, they become analytically tractable when conditioned to hypotheses of topological-level loop closure: if we factorize (3) through the definition of conditional probability:

$$p\left(s^t, m \middle| u^t, o^t\right) = p\left(s^t \middle| u^t, o^t\right) p\left(m \middle| s^t, u^t, o^t\right) \tag{4}$$

the map m becomes analytically tractable given the hybrid robot path s^t ([6],[18]). Next, we take advantage of having defined the map as a set of areas. Ideally, we could expect this partitioning to achieve the independence of observations within different areas, as illustrated in the DBN of Fig. 1. Mathematically, from this follows the conditional independence between robot path portions from different partitions, conditioned to the starting pose of the robot into each sub-map:

$$I\left({}^{i}s_k, {}^{j}s_l \middle| \left\{{}^{i}s_0, {}^{j}s_0\right\}\right) \quad \forall i \neq j,\ \ k,l > 0 \tag{5}$$

Unfortunately, in practice, partitioning the map will rarely generate strictly independent observations between sub-maps. However, we will still assume that (5) holds as an approximation while only a small quantity of information is lost. This enables the following factorization of (4):

$$\begin{aligned} & p\left(s^t, m \middle| u^t, o^t\right) \cong \\ & \prod_{k=1}^{n} p\left({}^{k}s^t \middle| u^t, o^t\right) p\left(m \middle| {}^{k}s^t, u^t, o^t\right) = \\ & \prod_{k=1}^{n} p\left({}^{k}s^t \middle| u^t, o^t\right) \prod_{k=1}^{n} p\left({}^{k}\mathcal{M} \middle| {}^{k}s^t, u^t, o^t\right) p\left(\left\{{}^{ab}\Delta\right\} \middle| \left\{{}^{k}\mathcal{M}, {}^{k}s^t\right\}, u^t, o^t\right) \end{aligned} \tag{6}$$

This theoretical result suggests that HTM-SLAM can be achieved through a set of n separate estimation processes, one for the robot path within each area. Then, map hypotheses are computed according to the corresponding estimated paths. Since the robot state s_t is a Markov process, we can sequentially estimate it via the Bayes rule:

$$\begin{aligned} & \overbrace{p\left(s_t \middle| u^t, o^t\right)}^{\text{Posterior}} \overset{Bayes}{\propto} \\ & \overbrace{p\left(o_t \middle| u^t, o^{t-1}\right)}^{\text{Observation likelihood}} \int \underbrace{p\left(s_t \middle| s_{t-1}, o^{t-1}, u^{t-1}\right)}_{\text{Transition model}} \underbrace{p\left(s_{t-1} \middle| u^{t-1}, o^{t-1}\right)}_{\text{Prior}} ds_{t-1} \end{aligned} \tag{7}$$

Certainly, an exact solution for the integral above is not available in the general case, forcing an approximate approach. Furthermore, the intractable growth in the number of possible topological paths [22] imposes a sample based approximation. Notice that the estimated path s^t defines the topological structure of the HMT map, as illustrated with some examples in Fig. 2. In this work we compute (7) by means of a Sequential Importance Resampling (SIR) filter [6].

Assume that the distribution of the robot path until the last time step (t-1) is available as a set of M particles:

$$\left\{s^{t-1,[i]}\right\}_{i=1..M} \sim p\left(s^{t-1} \middle| u^{t-1}, o^{t-1}\right) \tag{8}$$

The filter in (7) can be implemented as a prediction step followed by the corresponding particle weights update:

$$\left\{s_t^{[i]}\right\}_{i=1..M} \sim p\left(s_t \middle| u^t, o^{t-1}\right) \quad , \omega_t^{[i]} \propto \omega_{t-1}^{[i]} p\left(o_t \middle| s_t^{[i]}, m^{[i]}\right) \tag{9}$$

Next, a selective resampling step is required to alleviate the particle depletion problem [1]. Recalling the Markov assumption, we can now extend the robot path estimation as:

$$s^{t,[i]} \triangleq \left\{s^{t-1,[i]}, s_t^{[i]}\right\} \tag{10}$$

and repeat all the steps in the SIR filter recursively over time. Remember that Rao-Blackwellization grounds on the map distribution to be analytically computable from (10), thus each particle carries its own map estimation that is updated independently. In our context this means that each particle carries a hypothesis of the current area's metric map only, since previous local sub-maps are removed from the history of particles each time the robot enters a new area. It may be sometimes desirable to compute the discrete probability mass function (PMF) of the topological path γ^t, e.g. for visualizing existing topological structure hypotheses. This operation can be simply achieved by summation over

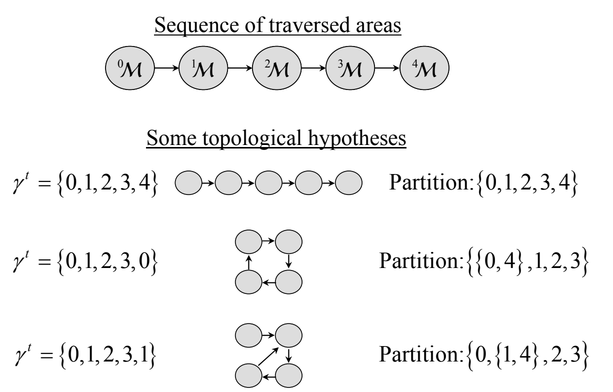

Fig. 2 Our approach takes the sequence of areas traversed by the robot (on the top), and estimates the topological structure of the environment according to the feasibility of possible partitions in that sequence. The bottom graphs show some examples of potential topologies and associated topological paths γ^t.

all the HMT path hypotheses:

$$P\left(\gamma^t = \gamma' \middle| u^t, o^t\right) = \iint p\left(x^t, \gamma', m \middle| u^t, o^t\right) dx^t dm \overset{\text{Particle aproximation}}{\cong} \sum_{i \in \Omega} \omega_t^{[i]} \qquad , \Omega = \left\{ i : \gamma^{t,[i]} = \gamma' \right\} \tag{11}$$

It is worth mentioning that virtually all the complexity in our approach rests on the drawing process from the distribution (9), mainly due to the estimation of the topological path. An implementation is proposed in a later section to deal with the complexity of this task in real-time applications, by means of dividing it into small, parallel sub-processes, and by postponing some operations.

IV. Relevant Elements Of HMT-SLAM

After exposing the theoretical foundations of our Bayesian approach to hybrid mapping, in this section we further clarify some key elements from this framework.

A. Partitioning the map

The hybrid map m is an annotated graph defined by the 2-tuple in (2), where nodes are the set of areas $\{{}^1\mathcal{M},\ldots,{}^n\mathcal{M}\}$ and edges represent coordinates transformation between them. Let ${}^{ab}\Delta$ be the transformation of coordinates from the area ${}^a\mathcal{M}$ into the reference system of ${}^b\mathcal{M}$. Provided for convenience that the first pose within each sub-map is the local coordinate reference, and being $\ominus$ the inverse pose composition operator, we have:

$$p\left({}^{ab}\Delta \middle| u^t, o^t\right) \triangleq p\left({}^a x_0 \ominus {}^b x_0 \middle| u^t, o^t\right) \tag{12}$$

In this work we compute this distribution by marginalization over all the particles:

$$\begin{aligned} & p\left({}^a x_0 \ominus {}^b x_0 \middle| u^t, o^t\right) \\ & = \iint p\left({}^a x_0 \ominus {}^b x_0 \middle| s^t, m, u^t, o^t\right) p\left(s^t, m \middle| u^t, o^t\right) ds^t dm \\ & \overset{\text{Particle aproximation}}{\cong} \sum_i \delta\left({}^a x_0 \ominus {}^b x_0 - {}^a x_0^{[i]} \ominus {}^b x_0^{[i]}\right) \omega_t^{[i]} \end{aligned} \tag{13}$$

which is a point mass approximation directly available from the RBPF employed for local metric SLAM.

The proposed model of sub-maps and coordinate transformations between them has been reported in other works ([7],[13],[14]), though we also propose a well grounded method that minimizes the loss of information in the map partitioning process: the minimum normalized-cut (min-Ncut) in the graph of observations [2],[30]. This is in contrast with previous works, which typically assume heuristic criteria.

To compute this partitioning, a weighted undirected graph is built where nodes are the robot observations and edge weights represent the covisibility between them. Periodically it is applied a highly-efficient spectral bisection algorithm to obtain the clustering that minimizes a given measure of independence (namely, the normalized cut). Notice that this bisection should be accepted only if the measurement rises over a given threshold which states the required independence between sub-maps. The partitioning process was already reported in detail elsewhere [2].

B. Uncertainty dereference/projection

An issue in any metric-topological approach to SLAM is the *projection* of uncertainty through different coordinate references. This process takes any pose distribution $p({}^a x)$ referenced to a given area ${}^a\mathcal{M}$, and computes its distribution $p({}^b x)$ relative to a second reference sub-map ${}^b\mathcal{M}$. The inverse operation is uncertainty *dereferencing*, which is a fundamental mechanism in our framework: each time the robot enters a new area, its pose uncertainty is dereferenced into the new coordinate system, as graphically illustrated in Fig. 3. This involves a pose inverse compounding operation:

$${}^b x = {}^a x \ominus {}^{ab}\Delta \tag{14}$$

Referencing uncertainty to the local frame provides a major advance over non-hybrid approaches to SLAM: the farther the robot is from the coordinate reference, the more complicated it becomes to model the pose probability distribution. For EKF methods this entails a raise of linearization and non-Gaussianity errors, whereas in RBPFs more particles are required to assure bounded approximation errors. Actually, these problems are a consequence of aiming to estimate global coordinates, which is avoided here.

The opposite situation (projecting a pose distribution towards a different area across the topological graph) is carried out by the pose composition operator $\oplus$:

$${}^a x = {}^{ab}\Delta \oplus {}^b x \tag{15}$$

which should be extended recursively if there are multiple transformations (edges) between the origin and target areas. Notice that, in general, many possible topological paths may exist between a given pair of areas, thus many different and probably inconsistent transformations are simultaneously applicable. We have adopted the solution reported in [3], where the Dijkstra algorithm is applied for finding the "shortest" topological path, in terms of smallest uncertainty.

This uncertainty projection has a direct application in finding absolute coordinates, a required step in computing a global map. Although our mapping approach does not need it at all, computing a global map relative to an arbitrary reference still is an appealing way of visualizing maps readable to humans, hence the interest in their construction. Experimental results in next sections show examples of such global maps, where inconsistencies may exist due to the problem mentioned above. One can devise optimization methods to reduce such inconsistencies [3] by means of ideas previously applied to consistent scan matching [17], but they are not applied here.

In our current implementation, uncertainty propagation in (14)-(15) is performed by Monte Carlo simulations, since distributions are already given as discrete samples from the RBPF. In the case of EKF-based SLAM, approximate closed-form solutions exist [7].

V. Implementation Framework

In section III we introduced the theoretic foundations of HMT-SLAM. Next a practical framework is presented which implements those ideas while keeping in mind that a mobile robot may demand accurate metric localization in real-time (e.g. for navigation or manipulation purposes), whereas maintaining the consistency of the topological map,

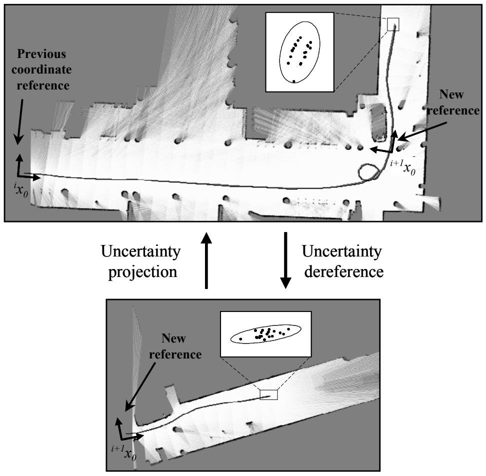

Fig. 3 Example of uncertainty dereference/projection from a real experiment. The top figure shows the evolution of particles as the set of estimated paths from each hypothesis. When entering into a new area is detected, the uncertainty is dereferenced into a new coordinate reference (see top figure). Relatively to the new area, the path uncertainty has been greatly reduced, as shown in the bottom figure. The ellipses indicate 99.97% confidence intervals for approximate Gaussian distributions.

thus it, solving loop closures, may be dealt in an any-time fashion.

The system is sketched in Fig. 4, where it can be seen the layered structure: metric local SLAM is performed in the low level, while more abstract (topological) representations are managed in upper levels. The inputs to the system are actions and observations from the robot, which are kept in a time-stamp-ordered queue until they can be processed. Within the system, there are a number of processes running concurrently which interact by reading and writing operations into the three levels, illustrated in Fig. 4: the local metric map of the current area, the sequence of traversed areas, and the space of topological path hypotheses. It must be remarked the parallel nature of the system, since the processes do not run in a predefined, sequential order. Next we describe the concurrent processes in the system and their relations with the different level described above:

- *Metric SLAM*: This process handles the robot localization and mapping within the local metric map for the current area, by processing actions and observations and integrating them into the Bayes filter. RBPFs represent an attractive choice here due to their constant-time operation, an essential feature to enable reliable real time low-level SLAM. This process is related to the estimation of the metric part of the robot path described in (7)-(10).

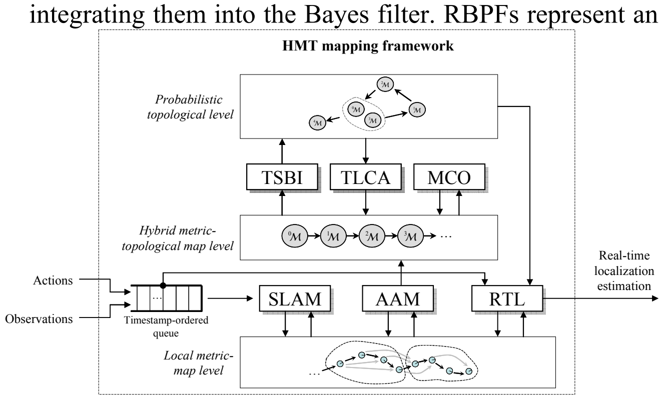

Fig. 4 Overview of a proposed implementation for our hybrid mapping framework, designed as a practical solution to the theoretical estimation process introduced in the paper. Please refer to the text for a detailed description.

- *Area Abstraction Mechanism* (AAM): Appearance-based methods are applied here to detect clusters of (approximately) independent observations in the sequence gathered by the robot ([2],[30]), i.e. whether the robot has entered into a new area. In such an event, observations from the last area are abstracted into a new area in the upper level. In our current implementation we assume that the partitioning of areas does not change with time, which does not represent a hurdle since the partitioning method is robust enough to produce approximately the same clusters of a given area independently of the robot path across it.
- *Topological Space Bayesian Inference* (TSBI): This process assigns values to the topological path γ^t of the robot according to the current map hypotheses [22], and is related to the topological part of the drawing process in (9). In our current implementation the inference on the topological path γ^t is postponed until the AAM algorithm starts a new sub-map. We have found that by doing so, the TSBI process can be performed through a simple maximum likelihood estimation (MLE) approach:

$$\gamma_t^{[i],*} = \arg\max_{\gamma_t} P\left(\gamma_t \left| u^t, o^{t-1}, s^{t-1,[i]}, m^{[i]}\right.\right)$$
$$P\left(\gamma_t \left| u^t, o^{t-1}, s^{t-1,[i]}, m^{[i]}\right.\right) \overset{Bayes}{\propto} P\left(\gamma_t\right) \prod_k p\left(o_k \left| \gamma_t, s^{t-1,[i]}, m^{[i]}\right.\right) \quad (16)$$

  where no information about the prior $P(\gamma_t)$ is assumed, i.e. we consider a uniform distribution over all paths. The method above simplifies the implementation, since (16) can be easily evaluated pointwise, while still performing well for large-scale map building. However, this simplification is not suitable for global HMT localization, an issue not addressed in the present work.
- *Topological Loop Closure Acceptance* (TLCA): Due to computation and storage limitations, it may be required that the robot forgets part of its topological path γ^t. The TLCA process performs this task, which is equivalent to accepting part of the topological structure hypothesis as correct. However, in practice this can be done for highly dominant (or unique) hypotheses, which can be determined by evaluating (11), hence the loss of information shall be negligible in most cases.
- *Maps Consistency Optimization* (MCO): This task is in charge of optimizing the relative pose of observations within each sub-map. The method from Lu and Milios [17], running in $O(n^2)$ with the map size, is appropriate here since the size of the local maps is bounded.
- *Real Time Localization* (RTL): This process guarantees an estimation of the robot position in a timely fashion. If the input queue is empty, the best pose estimation is s_t, already updated by the SLAM algorithm. However, if there is pending actions in the queue, the RTL computes the *prior* distribution, e.g. $p(s_{t+1}| s_t, u_{t+1})$, as a more updated estimation of the actual robot pose. This is clearly not the optimal solution, but it can be easily updated in real time.

Although our implementation does not exploit all the potential of the developed theoretical framework (mostly due to the postponed topological loop-closure through maximum likelihood estimation), it has demonstrated its suitability for large-scale map building of complex environments.

VI. Experimental Validation

We have tested our mapping framework with two different data sets, both with odometry readings and laser range scans in large-scale planar scenarios. One data set was gathered by the authors at the university of Malaga, and comprises almost 5000 laser scans collected along a 1.9Km path. The other data set was recorded at Edmonton Convention Centre by Nick Roy, and it is freely available online. Please, consider viewing the videos of the experiments [29] to get a better grasp of the results.

To compare the efficiency of our approach with previous methods we have also built the corresponding global metric maps with a highly efficient RBPF technique proposed in [10] for global metric mapping. The performance in both computation time and memory requirements is summarized in Table I. Results are for a 2.0GHz Pentium M (1Gb RAM), and for occupancy grid maps with a cell size of 12cm.

TABLE I
Performance Comparison Between Global RBPF-Mapping And Our Approach For HMT-SLAM

	Memory requirements		**Computation time**	
Method / Data set	**Global RBPF**	**HMT-SLAM**	**Global RBPF**	**HMT-SLAM**
Edmonton	84.22 Mb	27.97 Mb	39 min.	7 min.
Malaga	196.58 Mb	35.88 Mb	103 min.	24 min.

It is noticeable that HMT-SLAM outperforms global RBPF for both data sets. The improvement in the memory requirements follows from the fact that particles in our approach carry a hypothesis of the local metric map only, whereas in RBPF those hypotheses are for the whole global map. Therefore, the advantage of HMT maps becomes more and more relevant for increasingly larger environments. Regarding the lower computation time of our approach, it is a direct consequence of the reduced number of particles. However, by using a few particles only in local SLAM (we use 15 particles), our approach can attain a much better representation of the uncertainty (through Monte Carlo simulations of the uncertainty projection in (15)) than the one attainable by a global RBPF with a practical number of particles (e.g. less than 50). This turns into more precise loop closures in HMT-SLAM than in metric RBPF.

The map estimation before and after a loop closure are shown in Fig. 5(a)-(b) for the Malaga data set. In this situation, the robot enters a new area, labeled '11', leaving the previous area '10'. Then, the dominant topological path hypothesis becomes that of establishing the partition {'1','11'}, that is, the most likely explanation for robot observations is that area '11' actually corresponds with the area '1'. It can be appreciated in Fig. 5(c)-(d) how the loop closure affects the pose uncertainty of surrounding areas due to the introduction of a new edge that modifies the Dijkstra shortest paths employed to generate those global maps (notice that ellipses in the figure exaggerate the actual uncertainty by a factor of 5 for ease of visualization). The final HMT maps built from the Malaga and Edmonton data sets are plotted in Fig. 5(e)-(f), respectively, as global maps where global coordinate references have been arbitrarily set to the first nodes in each map. Although GPS readings are not available to measure absolute localization errors, our work is not focused on obtaining an accurate global metric map, but on reliably estimating the topological structure of the environment, which has been successfully carried out in both experiments.

VII. Conclusions And Future Work

In this work we have introduced a new viewpoint for solving the problem of large-scale SLAM which consists of estimating the hybrid metric-topological (HMT) path followed by the robot. It has been demonstrated that our approximation is supported by the probabilistic structure of the SLAM problem under weak assumptions.

We have also presented a relatively simple implementation of our ideas in the form of a real-time/any-time system. This implementation has demonstrated to be efficient for mapping large scale environments with multiple loops, but further research is needed to fully exploit the versatility of the HMT-SLAM framework against harder problems. For example, mapping highly ambiguous environments, or efficiently solving the robot awakening problem within large (even partially unknown) environments, are issues that can be hardly dealt with existing methods. We believe that the proposed paradigm of HMT-SLAM is a promising approach for these problems.

References

[1] M.S. Arulampalam, S. Maskell, N. Gordon, T. Clapp, "A Tutorial on Particle Filters for Online Nonlinear/Non-Gaussian Bayesian Tracking", in *IEEE Transactions on Signal Processing*, v.50, pp. 174-188, 2002.

[2] J.L. Blanco, J. González, J.A. Fernández-Madrigal, "Consistent Observation Grouping for Generating Metric-Topological Maps that Improves Robot Localization", in *ICRA06*, pp. 818-823, 2006.

[3] M. Bosse, P. Newman, J. Leonard, M. Soika, W. Feiten, S. Teller, "An Atlas framework for scalable mapping", in *ICRA03*, pp.1899-1906, 2003.

[4] H. Choset, K. Nagatani, "Topological simultaneous localization and mapping (SLAM): toward exact localization without explicit localization", in *IEEE Trans. on Rob. and Auto.*, v.17, pp.125-137, 2001.

[5] M. Dissanayake, P. Newman, S. Clark, H.F. Durrant-Whyte, M.A. Csorba, "A solution to the simultaneous localization and map building (SLAM) problem", in *IEEE Transactions on Robotics and Automation*, v.17, no. 3, pp. 229-241, 2001.

[6] A. Doucet, N.D. Freitas, K. Murphy, S. Russell, "Rao-Blackwellised Particle Filtering for Dynamic Bayesian Networks", in *Proc. of the 16th Conf. on Uncertainty in Artificial Intell*igence, 2000.

[7] C. Estrada, J. Neira, J.D. Tardós, "Hierarchical SLAM: Real-Time Achúrate Mapping of Large Environments", in *IEEE Transactions on Robotics*, v.21, no. 4, pp. 588-596.

[8] D. Fox, "KLD-sampling: Adaptive particle filters and mobile robot localization", in *Advances in NIPS*, 2001.

[9] U. Frese, "A Discussion of Simultaneous Localization and Mapping", in *Autonomous Robots*, v.20, pp.25-42, 2006.

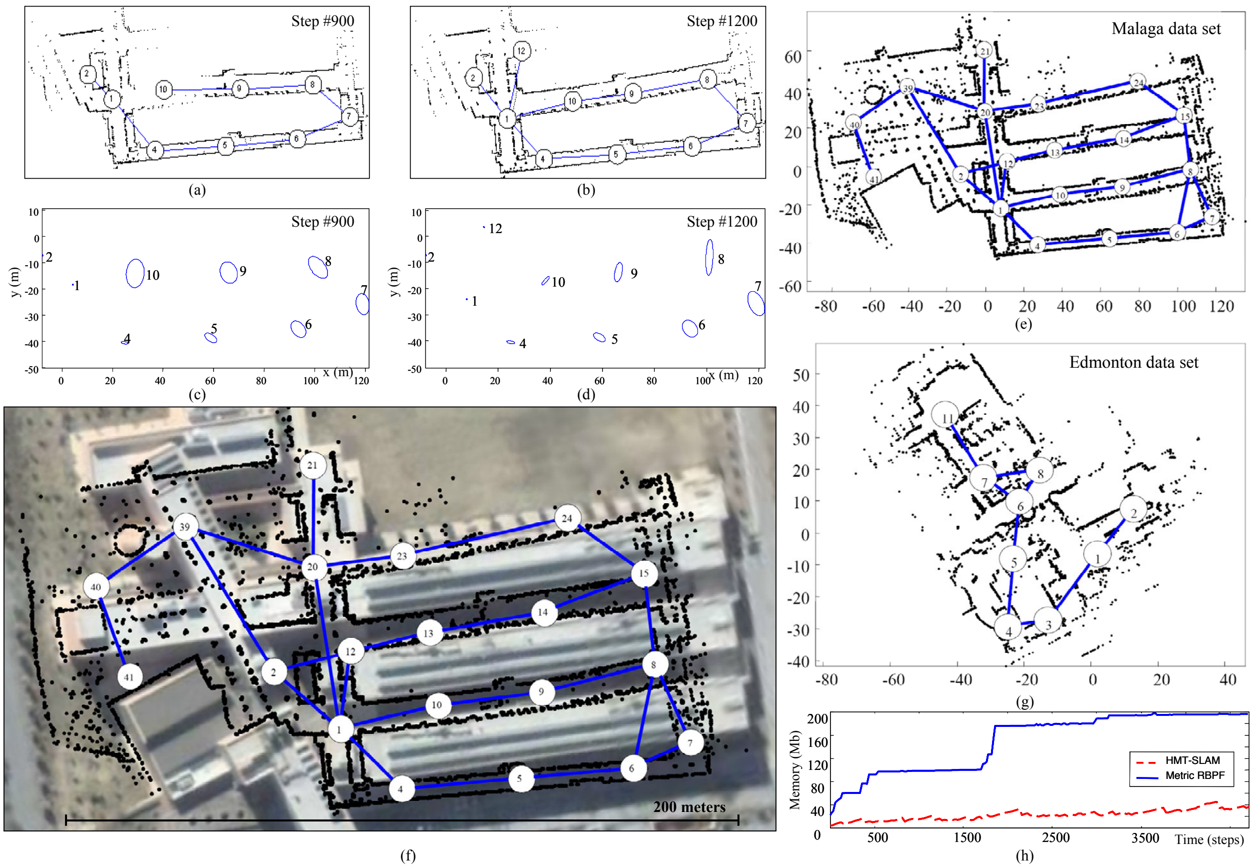

Fig. 5 (a)-(b) The map just before and after closing a loop by merging nodes {1, 11}. It is shown in (c)-(d) how this topological loop closure reduces the uncertainty in the position of surrounding nodes (uncertainty is represented by 95% confidence intervals, where uncertainty has been exaggerated by a factor of 5 for clarity). The final global map obtained for the Malaga data set is shown in (e), and in (f) overlapped with a satellite photo of the actual place. The graph (g) shows the global map obtained with our method for the Edmonton data set. The memory requirements of our approach in contrast with those of a global metric RBPF-mapping are also plotted in (h) for the Malaga data set.

[10]G. Grisetti, C. Stachniss, W. Burgard, "Improving Grid-based SLAM with Rao-Blackwellized Particle Filters by Adaptive Proposals and Selective Resampling", in *ICRA05*, pp. 2432-2437, 2005.

[11]G. Grisetti, G.D. Tipaldi, C. Stachniss, W. Burgard, D. Nardi, "Speeding-Up Rao-Blackwellized SLAM", in *ICRA06*, pp.442-447, 2006.

[12]J.S. Gutmann, K. Konolige, "Incremental mapping of large cyclic environments", in *IEEE International Symposium on Computational Intelligence in Robotics and Automation*, pp. 318-325, 1999.

[13]M.E. Jefferies, M.C. Cosgrove, J.T. Baker, W.K. Yeap, "The Correspondence Problem in Topological Metric Mapping-Using Absolute Metric Maps to Close Cycles", in *Int. Conf. on Knowledge-based Intelligent Information and Engineering Systems*, 2004.

[14]K. Kouzoubov, D. Austin, "Hybrid topological/metric approach to SLAM", in *ICRA04*, v.1, pp. 872-877, 2004.

[15]B. Kuipers, Y.T. Byun, "A Robot Exploration and Mapping Strategy Based on a Semantic Hierarchy of Spatial Representations", in *Robotics and Autonomous Systems*, v.8, pp. 47-63, 2001.

[16]B. Lisien, D. Morales, D. Silver, G. Kantor, I. Rekleitis, H. Choset, "The hierarchical atlas", in *IEEE Trans. on Robotics*, v.21, pp. 473-481, 2005.

[17]F. Lu, E. Milios, "Globally Consistent Range Scan Alignment for Environment Mapping", in *Autonomous Robots*, v. 4, pp. 333-349, 1997.

[18]M. Montemerlo, S. Thrun, D. Koller, B. Wegbreit, "FastSLAM: A factored solution to the simultaneous localization and mapping problem", *in Proc. of the AAAI Natl. Conf. on Art. Intell.*, pp. 593-598, 2002.

[19]P. Moutarlier,R. Chatila, "Stochastic multisensory data fusion for mobile robot location and environment modeling", in *5th Int. Symposium on Robotics Research*, pp. 207-216, 1989.

[20]H. Moravec, A. Elfes, "High resolution maps from wide angle sonar", in *Robotics and Automation*, v.2, pp. 116-121, 1985.

[21]J. Neira, J.D. Tardós, "Data Association in Stochastic Mapping Using the Joint Compatibility Test", in *IEEE Transactions on Robotics and Automation*, v.17, no. 6, pp. 890-897, 2001.

[22]A. Ranganathan, E. Menegatti, F. Dellaert, "Bayesian Inference in the Space of Topological Maps", in *IEEE Trans. on Rob.*, pp. 92-107, 2006.

[23]F. Savelli, B. Kuipers, "Loop-closing and planarity in topological map-building", in *IROS04*, v.2, pp. 1511-1517, 2004.

[24]R. Sim, P. Elinas, M. Griffin, A. Shyr, J.J. Little, "Design and analysis of a framework for real-time vision-based SLAM using Rao-Blackwellised particle filters", in *Canadian Conf. on Computer and Robot Vision*, 2006.

[25]C. Stachniss, G. Grisetti, W. Burgard, "Recovering Particle Diversity in a Rao-Blackwellized Particle Filter for SLAM After Actively Closing Loops", in *ICRA05*, pp. 655-660, 2005.

[26]J.D. Tardos, J. Neira, P.M. Newman, J.J. Leonard, "Robust Mapping and Localization in Indoor Environments Using Sonar Data", in *The International Journal of Robotics Research*, v.21, pp. 311-330, 2002.

[27]S. Thrun, "Robotic Mapping: A Survey", in *Exploring Artificial Intelligence in the New Millennium*, 2002.

[28]S. Thrun, A. Bücken, "Integrating grid-based and topological maps for mobile robot navigation", in *Nat. Conf. on Art. Intell.,* pp. 994-950, 1996.

[29]Website for multimedia content for experimental results, available in `http://www.isa.uma.es/C13/jlblanco.`

[30]Z. Zivkovic, B. Bakker, B. Krose, "Hierarchical map building using visual landmarks and geometric constraints", in *IROS05*, pp. 2480-2485, 2005.