Towards a Unified Bayesian Approach to Hybrid Metric-Topological SLAM

Jose-Luis Blanco, Juan-Antonio Fernández-Madrigal, and Javier Gonzalez

***Abstract*— This article introduces a new approach to Simultaneous Localization and Mapping (SLAM) which pursues robustness and accuracy in large-scale environments. Like most successful works on SLAM, we use Bayesian filtering to provide a probabilistic estimation which can cope with uncertainty in the measurements, the robot pose, and the map. Our approach is based on the reconstruction of the robot path in a hybrid discrete-continuous state space, which naturally combines metric and topological maps. There are two fundamental characteristics that set this work apart from previous ones: (i) the use of a unified Bayesian inference approach both for the metrical and the topological parts of the problem; and (ii) the analytical formulation of belief distributions over hybrid maps, which allows us to maintain the spatial uncertainty in large spaces more accurately and efficiently than previous works. We also describe a practical implementation which aims for real-time operation. Our ideas have been validated by promising experimental results in large environments (up to 30.000 m^2, a 2Km robot path) with multiple nested loops, which could hardly be managed appropriately by other approaches.**

***Index Terms*— Bayesian filtering, hybrid metric-topological maps, loop closure, mobile robots, Rao-Blackwellized particle filters, SLAM, topological maps.**

I. Introduction

SIMULTANEOUS Localization and Mapping (SLAM) is one of the central problems in mobile robotics, since the effective introduction of autonomous robots into real-life applications will undoubtedly require their operation in environments unknown at design time. The common formulation of the SLAM problem consists of building some kind of world representation from the sequence of data gathered by the robot, assuming no prior information about the environment and simultaneously localizing the robot using that representation. Different kinds of representations, or *maps*, have been proposed in the robotics and the artificial intelligence literature, ranging from low-level metric maps, such as landmark maps [12], [53] and occupancy grids [18], to topological graphs which contain high-level qualitative information [10], [24], [31], even multi-hierarchies of successively higher-level abstractions [20]. While existing techniques allow building maps of relatively large areas, SLAM remains a largely unsolved problem in relation to high-level representations and long-term operation within large-scale environments.

Manuscript submitted to IEEE Transactions on Robotics by May 16,2007, regular paper. Corresponding author: Jose-Luis Blanco

This work has been partly supported by the Spanish Government under research contract DPI2005-01391.

The authors are with the Department of System Engineering and Automation, University of Málaga, 29071 Málaga, Spain.

email: {jlblanco,jafma,jgonzalez}@ctima.uma.es

After an intense research during the last decade, it is clear that the most successful methods for SLAM are those based on probabilistic Bayesian estimation, which can manage noisy measurements and uncertainty in the robot location, the map, and, for those based on particle filters, also in data association [12], [16], [39], [57], [60]. Therefore, our proposal is grounded on the success and accuracy of techniques for metric localization and mapping within small-sized scenarios. It has also been shown that large-scale environments can be divided into *areas* of convenient sizes where these techniques can be applied efficiently to produce consistent *local sub-maps* [7], [19]. In our approach, this division of space depends on the nature of sensors in such a way that each area contains portions of the environment that are very likely to be sensed by the robot simultaneously, whereas parts of different areas will be rarely or never observed at the same time. We employ for this purpose existing methods [5], [63] based on this criterion of simultaneous visibility (called *overlap* in [5]). Notice that this kind of area does not correspond to logical or semantic divisions as could be interpreted by a human [37], such as a corridor or a room, but is based on the robot's sensory apparatus.

Using this definition of area, we introduce the concept of the hybrid metric-topological (HMT) path, which comprises the sequence of areas the robot has traversed (topological part) and its pose within each of them (metric part). Then, by considering the posterior belief distribution of the whole HMT path we can obtain the probability distribution over all the potential *topological structures* of the environment, an issue not addressed before simultaneously to the estimation of the metric poses between, and within, the areas. The resulting probabilistic map, called HMT map, represents the topology of the environment with graphs whose nodes (areas) are annotated with metric sub-maps and whose arcs (connections between areas) are annotated with the coordinate transformations between the corresponding areas. By conditioning the belief distribution of the map to the knowledge of the HMT path, we can represent these relative coordinate transformations in closed form, avoiding by design some problems that appear in global mapping with particle filters [36], [54]. The avoidance of absolute coordinates has been repeatedly proposed in the literature due to the difficulty of appropriately representing the uncertainty of poses far away from a global coordinate origin [7], [19], [22], [56].

Our approach, called HMT-SLAM, supports metric sub-maps of either landmarks or occupancy grid-maps. In the context of occupancy grid mapping, it naturally provides a correction of the robot path after closing large loops without rebuilding any global metric map, which has been pointed out sometimes as a weakness of grid-based, large-scale mapping.

Our work is related to some existing methods for hybrid mapping, specially to Hierarchical SLAM [19], and to the *Atlas* framework [7]. The fundamental contributions of our proposal in the context of these and other previous works are:

- The introduction of probability distributions over both the metric *and* the topological part of the robot path. Previous works have considered either the robot metric [26], [39], [43] or topological [46] paths separately. Apart from the mathematical consistency of a unified Bayesian approach, this formulation supports multiple topological hypotheses, and can be factorized in such a way that allows the uncertainty of large maps to be maintained accurately.
- A statistically grounded principle for the separation of the map into sub-maps. In [19], new sub-maps are started if the previous ones reach a given number of landmarks, while in [7] this is performed whenever a measure of the localization performance degrades. We propose instead to generate sub-maps that minimize a given measure[1] of covisibility or overlap between groups of observations, which allows us to set a grounded statistical model for HMT-SLAM as a Bayesian inference problem. Also, this provides the robot with topological structures that do not depend on external engineered knowledge, but on its own sensory system.
- In contrast to previous works (such as [7]), in HMT-SLAM all the hybrid map hypotheses are treated equally, associating different metric sub-maps to a given area if there exist multiple hypotheses about the topological path followed by the robot. This implies that the metric pose of the robot may be distributed around multiple modes even for particles with the same topological position. Our approach does not impose limits to the variety of the inferred distributions, whose complexity will uniquely be determined by the ambiguity of the environment and the uncertainty introduced by closing long loops.

The rest of this paper is outlined as follows. In section II we examine in more detail previous works on SLAM. Next, we provide the probabilistic foundations for our approach, while some of its key elements are discussed in depth in section IV. A practical system that implements these ideas is presented in section V, and experimental results with real robots in large-scale scenarios, as well as comparisons to other approaches, are discussed in section VI. Finally, some conclusions and future work are outlined.

II. Related Work

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

Pure metric approaches aim at reconstructing a representation of the environment where the relevant information is the metrical arrangement and characteristics of the map elements.

[1]Ideally, the cross-covariances in the case of landmark maps.

TABLE I

Summary of the Notation Employed in the Text

Symbol	Meaning
m	The HMT map (an annotated graph).
${}^{a}\mathcal{M}$	The local metric map for the area a.
${}^{b}_{a}\Delta$	The coordinate origin of area b relative to that of area a.
s_t	The robot HMT pose at time step t.
u_t, o_t	The robot actions and hybrid observations at time step t.
s^t, u^t, o^t	The sequences of robot poses, actions, and observations for time steps 1 to t.
${}^{i}s_{t'}, {}^{i}u_{t'}, {}^{i}o_{t'}$	A convenient way of referencing the robot poses, actions, and observations grouped into the area i such as the first elements are given for $t'=0$.
${}^{i}s^t, {}^{i}u^t, {}^{i}o^t$	The sequences of all the corresponding variables up to time step t.
ψ_t, z_t	The area-dependant and metric observations, respectively.
γ_t, x_t	The topological and metric parts of s_t at time step t, respectively.
γ^t	The topological path of the robot up to time step t.
Υ_t	The set of all known areas at time step t.
$s_t^{[k]}$	The k'th particle at time step t for the robot HMT pose.
$\omega_t^{[k]}$	Importance weight of the k'th particle at time step t.

Popular metric representations are landmark maps [1], [12], [53], [56], occupancy grids [18], [26], [41], and raw range scans [28], [35] (please refer to [57] for a more detailed classification). Some advantages of metric maps are their direct relation with robotic sensors and their value for some tasks such as motion planning or obstacle avoidance. There exist non-probabilistic approaches for building metric maps [25], [28], [35], although most works rely on probabilistic representations of the robot pose and the map, where Bayesian filtering is used to estimate the corresponding probability distributions [12], [53]. The hardest problem in those methods is data association (that is, establishing correspondences among observations and the map) [44], a problem that aggravates when the robot closes a loop since the uncertainty in the robot pose and the map increases as the robot explores new areas, i.e. the hardness of finding the correct data association increases with the scale of the maps, and in turn, establishing wrong correspondences severely compromises the consistency of the estimated maps. In fact, during the last years the research on the consistency of EKF-like solutions for SLAM has gained interest, since the approximations introduced by linearizing and assuming Gaussian distributions are known to impose a maximum length for a loop to be correctly closed [2], [8]. The idea of partitioning the environment into a sequence of sub-maps has been proposed as an improvement [45], [56], which is one of the basic motivations of our method.

A popular approach to the above mentioned problems of metric mapping is to employ a Rao-Blackwellized Particle Filter (RBPF) [13], [40], [43]. If we denote the map as m, the robot path as $x_{1:t}$ (or x^t for compactness), and the robot actions and observations as u_t and z_t, respectively, we can observe the following factorization of the SLAM posterior:

$$p(x^t, m|u^t, z^t) = p(x^t|u^t, z^t)p(m|x^t, u^t, z^t) \quad (1)$$

which is exploited in a RBPF and holds for any form of the probability densities. This factorization shows that, if we estimate the robot path through a separate particle filter, the map m can be estimated from the individual contributions of the observations z_t since they are conditionally independent given x^t. This is supported by the structure of the variables involved in SLAM, as shown in Fig. 1, which has been exploited successfully to build large maps both with occupancy grids [26] and landmarks (FastSLAM [39]). However, the number of samples required in these particle filters for closing a loop increases with the length of the loop, which may eventually turn into a storage capacity limitation since each particle carries a hypothesis of the whole map. Another limitation of RBPF for large-scale mapping is the loss of particle diversity when closing nested loops. Strategies exist for alleviating these problems [54], [55], but the underlying hurdles are just postponed. Another drawback of global mapping with RBPF is that the loss of diversity after closing a long loop typically leads to the total loss of the robot path uncertainty. An enlightening discussion about the problems of standard particle filters for large-scale mapping was recently presented in [36].

Building a topological map is an attractive alternative to metric mapping due to, among other properties, the reduced storage requirements and the good integration with symbolic planning of complex tasks [9], [15], [50]. However, pure topological maps are not suitable to solve SLAM. Although Bayesian estimation has been reported recently for these maps in [46], it is assumed a discrete set of “distinctive” places within the environment which must be correctly detected whenever the robot passes close to them: the only relevant sensorial information is that of being close to a distinctive place or not. We believe that a variety of sensors (like laser scanners, or cameras) can provide the robot a more accurate pose estimation through local metric sub-maps than that obtained by a purely topological approach, where most of the sensory data is simply ignored.

So far, it seems an appealing approach that of considering *hybrid maps*[2], where topological nodes contain local metric information [7], [19], [30], [32], [38], [59], [61]. In particular, the Atlas framework [7] and the work on hierarchical SLAM by Estrada *et al.* [19] contain interesting similarities with the present paper: both reference uncertainty to local coordinate frames and represent maps as topological graphs with local metric sub-maps. However, in these previous works loop closure has been considered only under the metric point-of-view, i.e. by finding the global metric coordinates transformation compatible with the loop closure. In turn, considering the whole HMT path of the robot leads to important advantages.

Finally, we should also highlight that the meaning of the topological part of our HMT map strongly differs from that considered in many other works. In the literature we can find works that consider distinctive places as nodes [9], [31], [32], while others cut the map into disjoint areas [19], [59]. Instead, our approach is closer to those where topological nodes are the result of *abstracting robot observations* gathered at a given area [5], [63]. Therefore, the size of the areas is automatically determined by the nature of sensors, more concretely, by the overlap between observations [5]. As an example, exploring a single room with a narrow field-of-view camera may result in many different areas, whereas a wide-angle laser scanner might probably lead to only one area. Our areas have no other special semantics.

[2]These maps are sometimes referenced as *hierarchical maps* in the SLAM literature. We use here the alternative term “hybrid” to avoid confusion with pure topological graph representations that involve abstraction processes, which are also usually called “hierarchical” [20], [23].

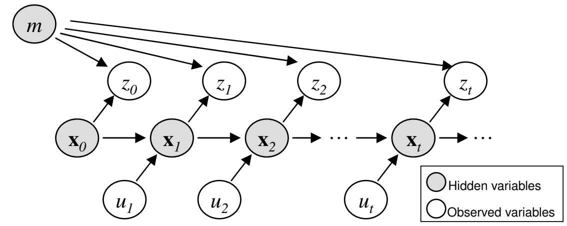

Fig. 1. The structure of the common formulation of SLAM as a Dynamic Bayesian Network. Actions u_t and observations z_t are the observed variables from which the robot path $x_{1:t}$ and the map m are estimated.

III. Probabilistic Foundations

The present proposal for HMT-SLAM is grounded on the sparsity of the relations between robot observations in a large environment: by our definition of *area*, observations within a given area will be highly related to one another, while observations belonging to different areas will not. This fact has been employed in a number of works to ease the construction of landmark maps [21], [34].

We start by defining a HMT map m of the environment as the annotated graph[3] that comprises the 2-tuple:

$$m = \langle \{ {}^i\mathcal{M} \}_{i \in \Upsilon_t}, \{ {}^b_a\Delta \}_{a,b \in \Upsilon_t} \rangle \tag{2}$$

Here the ${}^i\mathcal{M}$ are metric sub-maps for each area $i \in \Upsilon_t$, where Υ_t stands for the set of known areas at time step t. On the space of these hybrid maps, we will define beliefs as probability distributions. The local maps and the coordinate transformations ${}^b_a\Delta$ between the adjacent areas a and b are the annotations of the nodes and the arcs in the graph, respectively. Although this is the only information relevant for this work, it is worth mentioning that the graph could also maintain data about the kind each area is, the navigability between areas, or any other high-level knowledge useful in graph-based planning or symbolic reasoning [23]. We will consider conditional probability distributions over m given the information gathered by the robot up to some time step, where each particle may contain in turn different belief distributions for the metrical sub-maps and the coordinate transformations, and even a different number of nodes and arcs.

Accordingly, the robot is provided with a hybrid discrete-continuous description of its position within the map. Let $s_t = \langle \gamma_t, x_t \rangle$ denote the robot HMT pose at time step t,

[3]By “annotated graph” we mean a graph whose nodes and arcs can hold some non-topological information. This is an informal version of the term “typed attributed graph” in the graph transformation literature [17].

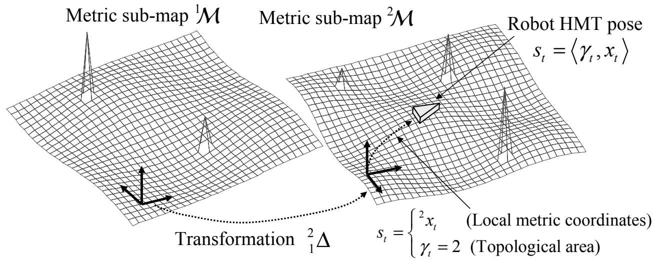

Fig. 2. Each metric sub-map $^i\mathcal{M}$ has its own local coordinate frame. The robot HMT pose is thus a discrete-continuous variable which states both the current area (the node in the topological part of the map) and the metric pose relative to the corresponding local frame.

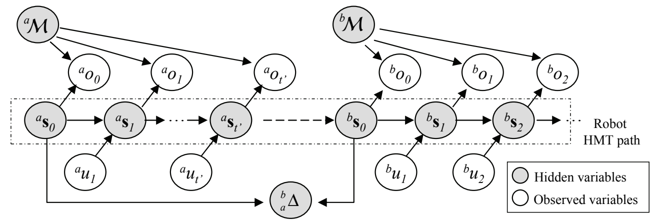

Fig. 3. The graphical model for HMT-SLAM. Here segments of the robot path are conditionally independent given the starting pose at each segment. The relative pose between areas is represented by $^b_a\Delta$, while the term $^i\mathcal{M}$ stands for the metric sub-map of area i.

where the discrete variable γ_t indexes the current area and the continuous variable x_t represents the metric pose relative to the local coordinate frame of that area, as represented schematically in Fig. 2.

We can now state the problem of HMT-SLAM as the estimation of the joint posterior of the robot HMT path s^t and the map m given the sequences of robot actions u^t and observations o^t up to time step t, that is:

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

For clarity, we denote sequences of variables over time with superscripts, i.e. $u^t = \{u_1, ..., u_t\}$. We also define o_t as containing the hybrid pair $\langle \psi_t, z_t \rangle$, with the purpose of conveniently setting metric observations z_t (such as range scans or landmarks extracted from images) apart from topologically dependant observations ψ_t (such as the recognition of a particular kind of area). This division renders especially useful when facing the problem of global localization.

In this paper we will approach HMT-SLAM under the point-of-view of sequential estimation, using the graphical model proposed in Fig. 3. For clarity, the first time step at each segment has been denoted as 0, and we add the current area as a left superscript to the name of variables, e.g. replacing s_t by $^{\gamma_t}s_{t'}$. We can observe the following interesting statistical properties in the structure of HMT-SLAM under the proposed sequential estimation model.

Proposition 1. Given the distribution of the robot HMT path, and due to the conditional independence between the local sub-maps, the joint posterior of all the sub-maps $i \in \Upsilon_t$ can be factorized as:

$$p(\{^i\mathcal{M}\}_{i\in\Upsilon_t}|s^t, u^t, o^t) = \prod_{i\in\Upsilon_t} p(^i\mathcal{M}|^i s^t, ^i o^t) \tag{4}$$

This property holds in general, regardless the choice of the metric map representation.

The conditional independence between the sub-maps given the robot path can be clearly seen in Fig. 3, where the robot path *d*-separates [49], [62] all the possible paths between any pair of sub-maps $^i\mathcal{M}$.

Proposition 2. The metric sub-maps $^i\mathcal{M}$ and the coordinate transformations $^b_a\Delta$ are conditionally independent given the robot HMT path.

This property can be verified in Fig. 3 by noticing that the robot path also *d*-separates the set of metric sub-maps and the set of arcs $^b_a\Delta$.

Proposition 3. The posterior distribution of the robot HMT path s^t can be factorized into the product of its segments through the different areas i:

$$p(s^t|u^t, o^t) = \prod_{i\in\Upsilon_t} p(^i s^{t'}|^i u^{t'}, ^i o^{t'}) \tag{5}$$

The reason of this property is that consecutive robot poses grouped into different areas, like $^a s_{t'}$ and $^b s_0$ in Fig. 3, become independent variables in our model. This may seem surprising at first glance, since these poses will often represent close, consecutive positions along the robot path, and they are actually related by a robot action (e.g. odometry) and observations of roughly the same place. Thus, the assumption of independence between *pure metric poses* would certainly lead to a significant loss of information. The key point here is that we assume their independence as *hybrid robot poses* (metric plus topological), where the metric coordinates are referenced to different local frames. The information from the last robot action ($^a u_{t'}$ in Fig. 3) is not lost, but incorporated into the corresponding $^b_a\Delta$ variable.

Nevertheless, it must be noticed that the cross-covariance between robot poses at different areas is definitively lost in HMT-SLAM. In practice, partitioning the map will rarely generate strictly independent observations between local maps, which renders our approach as an approximate solution to SLAM. The loss of information, however, can be minimized as much as desired by using grounded methods in the process of deciding when to start a new area [5], [63]. As a limit situation, if we desire no loss of information at all, HMT-SLAM degenerates into the common global metric SLAM. For simplicity, we also assume here that the methods for defining areas give us roughly the same results independently of the concrete robot path. This is required for identifying two different areas as the same physical placement, a requirement for closing loops at the topological level. Although this may lead to sub-optimal partitions, our partitioning method will always generate the best ones in terms of the maximum independence between observations. Our results show that, under the common assumption of an almost static world, our method is appropriate for practical situations.

Based on the above statistical properties of HMT-SLAM, we propose the following solution to the problem of estimating the posterior of Eq. (3). We start by using the Rao-Blackwellization approach [13] to first factorize the joint posterior into one component for the robot HMT path s^t and another one for the map m:

$$p(s^t, m|u^t, o^t) = p(s^t|u^t, o^t)p(m|s^t, u^t, o^t) \tag{7}$$

In this way we reduce the dimensionality of the joint path-map space to that of the robot HMT path only (s^t), which can then be estimated using a particle filter. For each particle it is computed the analytical conditioned distribution of the map (m) associated to the corresponding path hypothesis. Writing down the analytical part of Eq. (7) and given the conditional independence between the elements of m (proposition 2 above), we get:

$$\begin{aligned} p(m|s^t, u^t, o^t) &= p(\{{}^i\mathcal{M}\}_{i\in\Upsilon_t}, \{{}^b_a\Delta\}_{a,b\in\Upsilon_t}|s^t, u^t, o^t) \\ &= p(\{{}^i\mathcal{M}\}_{i\in\Upsilon_t}|s^t, o^t)p(\{{}^b_a\Delta\}_{a,b\in\Upsilon_t}|s^t, u^t, o^t) \end{aligned} \tag{8}$$

As shown in proposition 1, the first distribution in this product can be factorized and each sub-map computed using closed form equations for landmark maps [39] and occupancy grids [42], [58]. The second element in Eq. (8) is the joint posterior of the variables ${}^b_a\Delta$. As typically assumed in the SLAM literature, we propose to use Gaussians to model these relative poses, that is:

$${}^b_a\Delta \sim \mathcal{N}({}^b_a\bar{\Delta}, {}^b_a\Sigma) \tag{9}$$

We can now perform a Bayesian fusion in closed-form for each arc ${}^b_a\Delta$ in the HMT map, simply by:

$$p({}^b_a\Delta|s^t, u^t, o^t) = \prod_{j=1}^{L} \mathcal{N}({}^b_a\bar{\Delta}'_j, {}^b_a\Sigma'_j) \tag{10}$$

where L is the number of times the robot has moved between the areas a and b. The parameters of the j'th Gaussian in Eq. (10) can be obtained from the metric path of the robot (when the robot explores new areas), or from map alignment procedures (when considering the hypothesis of a topological loop closure). This means that each time the robot moves between a given pair of areas, the estimation of their relative pose is refined and the uncertainty is reduced. Note as well that each variable ${}^b_a\Delta$ can be estimated independently as a consequence of proposition 3, which is consistent with the construction of a map of relative poses. Therefore,

$$p(\{{}^b_a\Delta\}|s^t, u^t, o^t) = \prod_{a,b} p({}^b_a\Delta|s^t, u^t, o^t) \tag{11}$$

Next we address the non-analytical estimation of the robot path, the first term in Eq. (7). We estimate this posterior sequentially by Bayesian filtering:

$$p(s^t|u^t, o^t) \propto p(o_t|s^t, u^t, o^{t-1})p(s^t|u^t, o^{t-1}) \tag{12}$$

Like in [46], a particle filter is employed as a convenient representation of the topological part of the robot path (γ^t). Therefore, if we assume a set of P weighted particles distributed approximately according to the posterior for time step $t-1$:

$$\{s^{t-1,[k]}\}_{k=1..P} \sim p(s^{t-1}|u^{t-1}, o^{t-1}) \tag{13}$$

we can sequentially generate the particles for the next time step t by drawing samples from a given proposal distribution $q(s_t|s^{t-1}, o^t, u^t)$ and updating the importance weights $\omega_t^{[k]}$ accordingly. We consider here the optimal proposal for each particle to be $p(s_t|s^{t-1,[k]}, o^t, u^t)$, which has been demonstrated to minimize the variance of the weights [14]. Having an exact expression for this proposal means that the generated particles will be distributed according to the true posterior. As shown in the derivation in Eq. (6), we can put the optimal proposal for our particle filter in a form that allows us sampling a new particle $s_t^{[k]}$ in two steps: to draw firstly the topological position $\gamma_t^{[k]}$ using a topological transition model (discussed

$$\begin{aligned}
s_t^{[k]} &= \left\langle \gamma_t^{[k]}, x_t^{[k]} \right\rangle \sim q\left(s_t|s^{t-1,[k]}, u^t, o^t\right) \\
&= p\left(s_t|s^{t-1,[k]}, u^t, o^t\right) \overset{\text{Bayes}}{\propto} p\left(s_t|s^{t-1,[k]}, u^t, o^{t-1}\right) p\left(o_t|s_t, s^{t-1,[k]}, u^t, o^{t-1}\right) \\
&= p\left(\gamma_t, x_t|s^{t-1,[k]}, u^t, o^{t-1}\right) \overbrace{p\left(z_t|s_t, s^{t-1,[k]}, u^t, o^{t-1}\right) p\left(\psi_t|\gamma_t, s^{t-1,[k]}, u^t, o^{t-1}\right)}^{\text{Independence between } z_t \text{ and } \psi_t} \\
&= \overbrace{P\left(\gamma_t|s^{t-1,[k]}, u^t, o^{t-1}\right) p\left(x_t|\gamma_t, s^{t-1,[k]}, u^t, o^{t-1}\right)}^{\text{Definition of conditional probability}} p\left(z_t|s_t, s^{t-1,[k]}, u^t, o^{t-1}\right) p\left(\psi_t|\gamma_t, s^{t-1,[k]}, u^t, o^{t-1}\right) \\
&\rightarrow \left\{ \begin{array}{ll} 1^{\text{st}} \text{ step:} & \gamma_t^{[k]} \sim P\left(\gamma_t|s^{t-1,[k]}, u^t, o^{t-1}\right) p\left(\psi_t|\gamma_t, s^{t-1,[k]}, u^t, o^{t-1}\right) \\ 2^{\text{nd}} \text{ step:} & x_t^{[k]} \sim p\left(x_t|\gamma_t^{[k]}, s^{t-1,[k]}, u^t, o^{t-1}\right) p\left(z_t|x_t, \gamma_t^{[k]}, s^{t-1,[k]}, u^t, o^{t-1}\right) \end{array} \right.
\end{aligned} \tag{6}$$

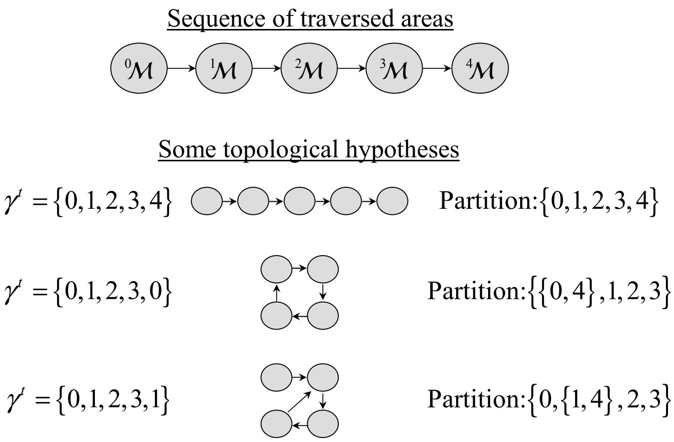

Fig. 4. Our approach takes the sequence of areas traversed by the robot (on the top), and estimates the topological structure of the environment by considering some of the potential rearrangements of the sequence. The bottom graphs show some examples of partitions (rearrangements), the associated topological paths γ^t, and the resulting map topologies.

in a later section), and to draw then the metric pose $x_t^{[k]}$ from the conditional distribution of the obtained topological position. This procedure is supported by our knowledge of the conditioned density of the metric pose given the topological position. The metric sample can be obtained from exact equations for landmark maps [40] or from approximations for occupancy grids [26]. This process is repeated for each time step, performing resampling [48] whenever the effective sample size [33] of the RBPF falls below a given threshold, e.g. the 50%. Drawing hypotheses for the topological position γ_t is arguably the most complex step in HMT-SLAM. Later on we discuss an implementation aiming at real-time execution which has given good results.

Observe that each hypothesis of the topological path γ^t implies a different topological structure of the environment (and thus, a different hypothesis for m) by means of clustering the sequence of areas traversed by the robot [46]. Since each particle may maintain a different HMT hypothesis we have a probability distribution over the possible topologies of the map, as illustrated with a few examples in Fig. 4.

IV. Relevant elements in HMT-SLAM

After discussing the theoretical foundations of our approach, in this section we deal with some practical issues that arise in HMT-SLAM.

A. The Transition Model of the Topological Position

A fundamental part of HMT-SLAM is the estimation of the topological path of the robot. In this work we have addressed this problem sequentially via particle filtering and considering the optimal proposal distribution for generating new hypotheses. As shown in Eq. (6), this leads to drawing samples $\gamma_t^{[k]}$ from the product of two terms: the transition model of the topological position $P\left(\gamma_t | s^{t-1,[k]}, u^t, o^{t-1}\right)$, denoted in the following as $P\left(\gamma_t | d^{[k]}\right)$ for clarity, and the appearance observation model $p\left(\psi_t | \gamma_t, s^{t-1,[k]}, u^t, o^{t-1}\right)$. For shortening this paper we will not detail the latter here, although

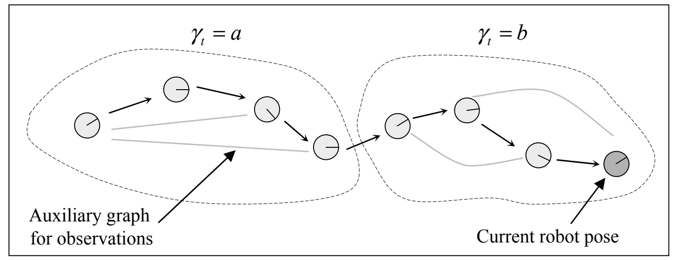

Fig. 5. An auxiliary graph of robot observations can be used to detect when the robot enters into a new area through graph bisection techniques.

some possible implementations have been reported by other authors [11].

Due to its discrete nature, the topological transition model assigns a probability to a reduced number of potential events:

- To stay at the same topological area (simply $\gamma_t^{[k]} = \gamma_{t-1}^{[k]}$).
- To enter into an unexplored area, $\gamma_t^{[k]} \notin \Upsilon_{t-1}$.
- To close a topological loop, that is, $\gamma_t^{[k]} \in \Upsilon_{t-1}$ with $\gamma_t^{[k]} \neq \gamma_{t-1}^{[k]}$.

The topological position of the robot γ_t is therefore a piecewise constant sequence over time t. For example, consider the robot path shown in Fig. 5, where the robot topological pose is a or b in the two separate parts of the path.

We can apply existing methods for partitioning a sequence of robot observations while minimizing some measure of similarity or overlapping between them (the exact measure to minimize may depend on the actual sensors and map representations [5], [63]). In short, these methods build an auxiliary graph whose nodes are the robot poses where observations were taken, and undirected weighted arcs between nodes keep the measure of similarity between the observations, as represented in Fig. 5. Then, the minimum normalized cut of the graph can be computed by means of spectral bisection [51]. Lower cut values mean weaker relation between the observations in the two subgraphs. A more detailed exposition of this process can be found elsewhere [5]. Thus, the robot has moved to a different area if the resulting cut is below a certain limit that settles the maximum allowed dependence between adjacent sub-maps. If such a partitioning happens, the immediate past of the robot path is updated to the new topological position. Note that this process can be performed only once for groups of particles that share the same topological path.

Going back to the topological transition model we can now specify that when the above partitioning method does not find a sufficiently independent partition of the auxiliary graph (that is, the robot has not entered a new area), we have $P\left(\gamma_t | d^{[k]}\right) = 1$ for $\gamma_t = \gamma_{t-1}^{[k]}$, and $P\left(\gamma_t | d^{[k]}\right) = 0$ otherwise. In contrast, when a new area is entered, the transition model considers the possibility of the robot having closed a loop by means of $P\left(\gamma | d^{[k]}\right) \propto F(\gamma)$, and the possibility of having entered into an unexplored area by $F(\gamma') = \eta$, where η is a constant and all the probabilities are scaled such as they sum up to unity. The function $F(\gamma)$ provides a measure of how likely it is to have arrived at the previously known area γ through a loop closure. This measure should incorporate the metric information of the arcs along the topological path

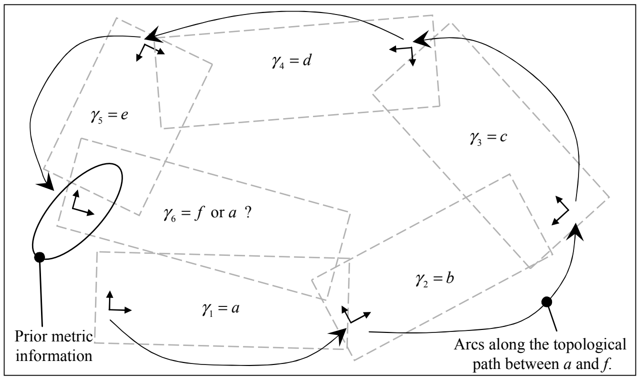

Fig. 6. Is the robot after the loop at area a or at f? The topological transition model assigns a probability to each of these possibilities by accounting for the robot observations and the prior metric information (and uncertainty) contained in the arcs of the HMT map.

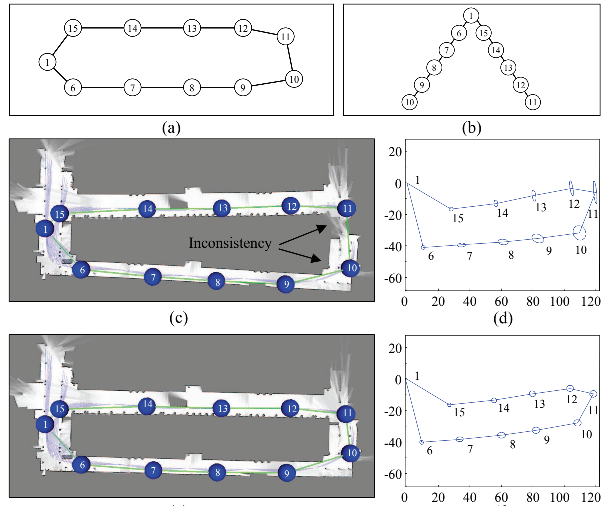

Fig. 7. (a) An example of a HMT map after closing a loop. (b) A tree representation for finding the shortest path from node 1 to any other node. (c)–(d) The global map and the uncertainty for each area obtained by using the shortest-path method for computing global coordinates. (e)–(f) The same map obtained by using a globally consistent method. The ellipses represent 95% confidence intervals magnified by factor 5 for clarity.

between the two nodes (refer to Fig. 6), although it can also rely on high-level topological reasoning [15], [47]. An interesting choice for this function is the likelihood of the last observations in relation to the metric sub-map of the candidate area γ. Since we would need the exact robot pose *after* the loop closure to compute this measure, we rather take the mathematical expectation using the prior distribution of the robot pose, given by the information contained in the arcs. In this way, we obtain candidates for the loop closure that both are within the scope of the a priori metric uncertainty and provide a good prediction of the last observations.

B. Probability Distribution over Topologies

Sometimes it may be desirable to compute the discrete probability mass function (PMF) of the topological path γ^t, for example for visualizing the topological structure hypotheses considered by the filter at a certain instant of time. This operation can be achieved by marginalization which for our particle filter simply becomes:

$$P(\gamma^t = \tilde{\gamma}^t | u^t, o^t) = \sum_{k \in \Omega} \omega_t^{[k]} \quad , \Omega = \{k : \gamma^{t,[k]} = \tilde{\gamma}^t\} \tag{14}$$

for each desired value of $\tilde{\gamma}^t$. Some examples of these distributions are shown in Fig. 10.

C. Obtaining Global Maps

During the normal operation of the robot using HMT-SLAM it will usually be enough to reference the robot coordinates to the coordinate frame of the current area. However, we could desire sometimes to compute absolute coordinates relative to some arbitrary reference, for example for constructing a global map for debugging or simple visualization. We describe next two possible methods for computing the absolute coordinates of all the areas of a map taking one of them as reference.

The first approach, proposed by Bosse *et al.* in [6] and adopted in our previous work [4], consists of applying the Dijkstra algorithm for finding the shortest topological path from the reference area to the rest. In this way, the topological part of a HMT map is transformed into a tree that encodes all these "optimal" paths, as shown with an example in Fig. 7(a)–(b). If we denote the reference area as a, and the area we are interested in as b, we will find the corresponding coordinate transformation ${}^b_a\Delta$ in the HMT map only if a and b are adjacent areas. Otherwise, this pose can be computed by sequentially composing pose changes along the shortest topological path $\{a, 1, 2, 3, ..., n, b\}$, that is:

$${}^b_a\Delta = {}^1_a\Delta \oplus {}^2_1\Delta \oplus ... \oplus {}^b_n\Delta \tag{15}$$

where $\oplus$ is the metric pose composition operator [52]. Given the conditional independence between all the relative poses along the topological path, a Monte Carlo simulation can be employed to generate samples of ${}^b_a\Delta$ from independent samples of all the elements in Eq. (15). The problem with this approach is that the existence of loops leads to inconsistent global coordinates. To illustrate the problem, please refer to Fig. 7(c)–(d) which represent the global coordinates and the associated uncertainties for each area of a real HMT map. The information in the arc between nodes 10 and 11 has not been considered by the shortest-path approach, hence we find an inconsistency in the resulting map between these two areas.

To solve this, we propose the alternative approach of applying methods for globally consistent poses estimation that have been reported in the past for networks of laser scans [35]. In this kind of method, we iteratively compute the approximate optimal poses of all the nodes by minimizing the linearized version of a cost function which includes all the constraints between adjacent areas. This algorithm typically converges in a few iterations, and the so obtained global maps are free of inconsistencies, as can be observed for the example in Fig. 7(e)–(f). This is the method employed for the rest of global maps that illustrate this article.

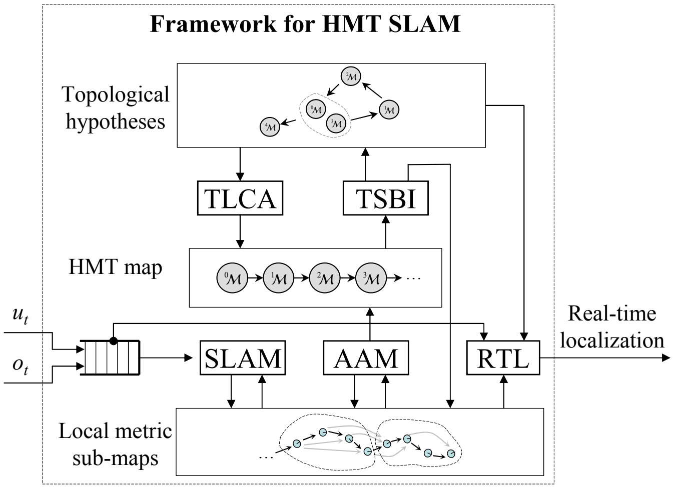

Fig. 8. Overview of the implementation proposed as a practical solution to HMT-SLAM. Here local metric SLAM is solved efficiently in real-time, while the topological structure of the map is estimated concurrently in a delayed fashion. In this way we prioritize the update of the metric position of the robot within the current area. Please refer to the text for further details.

D. Long-Term Operation

It is worth to highlight a crucial property of HMT-SLAM that makes it specially suitable for long-term operation of a mobile robot. The statement of HMT-SLAM as the estimation of the density in Eq. (3) includes the robot HMT path s^t, whose dimensionality always increases over time. In this sense, it may seem that our approach suffers from the same problem that global mapping with RBPFs, i.e. performing estimation into a state-space of unbounded dimensionality. However, hypotheses for the whole topological path γ^t can be forgotten until the point where the differences between the particles begin. That is, if at a given instant of time all the particles agree about the current topological structure, there is no need to maintain in memory several map topologies, with their corresponding metric local maps. Unlike the case of purely metric SLAM, this does not mean a loss of the estimated uncertainty along the robot path, since in HMT-SLAM this spatial uncertainty is maintained in a parameterized form by the conditional distributions for ${}^b_a\Delta$ (refer to Eq. (10)).

V. Implementation Framework

In sections III and IV we have introduced the theoretic formalization of HMT-SLAM. In the following we present a practical framework which implements those ideas. A relevant issue here is that a mobile robot may demand accurate metric localization in hard real-time (e.g. for navigation or manipulation purposes), while maintaining the consistency of the topological map (i.e. solving loop closures) can be performed in the “background” since reasonable delays are acceptable.

An overview of the proposed framework is presented in Fig. 8, where we can observe its layered structure. Metric local SLAM is performed at the low level, while topological representations are managed at the upper levels. Symbolic reasoning or task planning would fit in additional layers above these. The inputs to the system are actions u_t (typically produced by a planning level), and observations o_t (acquired by 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 means of read and write operations on the data held in the three levels represented in Fig. 8: 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, that is, the processes do not need to run in a predefined, sequential order. Next we describe these processes and their relations with the theory presented previously.

Metric SLAM: It handles the robot localization and mapping within the metric sub-map for the current area by processing actions and observations. Conceptually, this process is involved in estimating the metric part of Eq. (6). For occupancy grids, RBPFs have a complexity linear with the number of particles. Therefore, for a static number of particles we have a constant computation time independently of the size of the local area, which is a requisite if we desire a hard real-time estimation. In practice, we would need a variable number of particles depending on the number of topological path hypotheses considered at each instant of time, which may increase after closing long loops. If we want metric SLAM to have a bounded time complexity even in those cases we could impose a maximum number of particles in the filter, at the cost of disregarding the most unlikely topological hypotheses in the TSBI process.

Area Abstraction Mechanism (AAM): Grounded methods are applied here to detect clusters of (approximately) independent observations in the sequence of observations gathered by the robot [5], as already described in section IV-A.

Topological Space Bayesian Inference (TSBI): The topological transition model described in section IV-A is applied here whenever the AAM detects that the robot has moved into a new area. In our current implementation, the local metric-maps for the particles are not updated until TSBI gives us the hypotheses for the new topological position of the robot. Once the topological transition model has been evaluated, all the particles are updated to account for the changes, which include changing the coordinate references, removing part of the current metric sub-map in the case of entering into a new area, and in the case of a loop closure, loading the contents of the known area into the sub-map. Regarding the computational complexity of the TSBI, a straight-forward implementation exhibits quadratic complexity in the number of known areas $|\Upsilon_t|$. This complexity is imposed by the Dijkstra algorithm used in the computation of the a priori metric information (see IV-A).

Topological Loop Closure Acceptance (TLCA): As discussed above, it is unpractical to keep the whole history of the robot hybrid path for long-term operation. The TLCA process is in charge of accepting part of the topological structure as definitively correct. To avoid the risk of losing the valid (unknown) hypothesis in this process, it only operates when the system contains highly dominant (or unique) hypotheses for the robot topological path γ^t. In our current implementation this process also performs a fine alignment of the local grid maps before building the metric map for the topological node

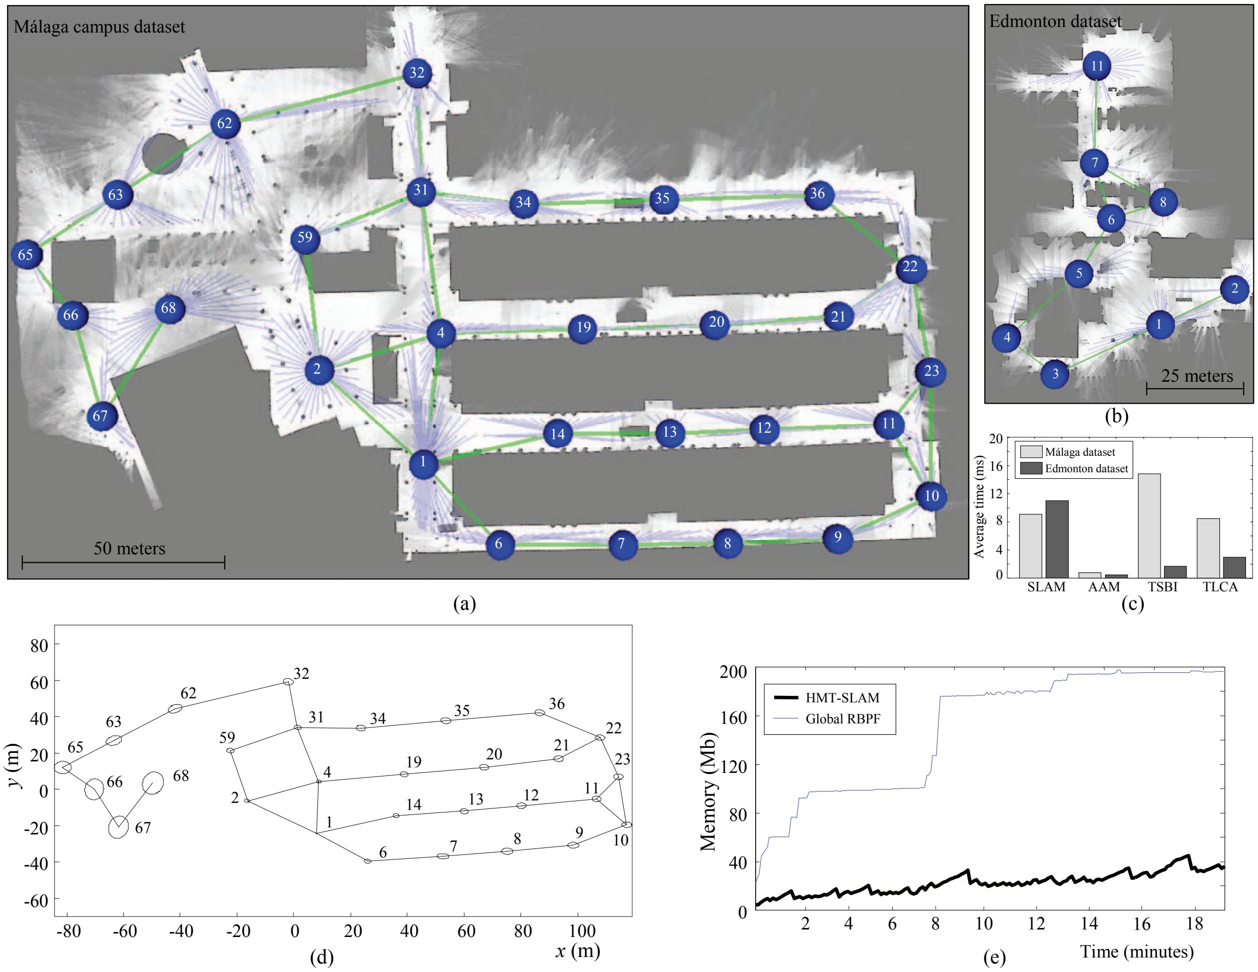

Fig. 9. (a)–(b) The HMT maps generated from the Málaga and Edmonton datasets, respectively, shown as global maps relative to the first topological area (labeled as '1'). (c) The average time (per action-observation pair) taken by each of the processes in our HTM-SLAM implementation. (d) The globally consistent poses of each area for the Málaga map. The ellipses represent 95% confidence intervals magnified by factor 5 for clarity. (e) Comparison of the memory requirements over time between our approach and metric mapping with RBPF, both for the Málaga map.

resulting from the fusion.

Real Time Localization (RTL): This process guarantees an estimation of the robot position in a timely fashion. If the input queue of actions and observations is empty, the best estimation of the HMT pose s_t has been already updated by our HMT-SLAM algorithm. On the contrary, when there are 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. We can easily compute this prior for the robot metric pose in real-time, for example, by accumulating odometry readings. The obtained pose estimations will not be the optimal ones, but as long as the metric SLAM process updates its estimation in a timely fashion, the estimate from RTL will be simultaneously accurate, and fast to obtain.

VI. Results

We have tested our implementation of HMT-SLAM with two different datasets, both comprised of odometry readings and laser range scans in large planar scenarios containing several loops. The first dataset [3] was gathered by the authors at the campus of the University of Málaga (Spain), and contains almost 5000 laser scans collected along a 2Km path. The second dataset was recorded at the Edmonton Convention Centre (Canada) by Nicholas Roy, and is freely available online [29].

For illustrating the experiments the resulting HMT maps for both datasets are shown in Fig. 9(a)–(b) as global maps (recall section IV-C), taking the first area in each map as the reference[4]. We also overlap the most likely topological structure inferred by our approach to visualize the existing nodes and arcs. In contrast to global metric mapping with RBPF, our approach is able to maintain the uncertainty in all the relative poses between the areas, as can be seen with the (globally consistent) uncertainties represented in Fig. 9(d). Recall that in RBPF-based global mapping, resampling steps eventually lead to a total loss of the represented uncertainty.

To compare the efficiency of other methods to our HMT-SLAM implementation, we have also built the corresponding maps with an efficient RBPF-based technique for (global)

[4]This paper has supplementary downloadable material available at http://ieeexplore.ieee.org, provided by the authors. This includes videos showing the mapping process for both datasets. This material is 48.1 MB in size.

TABLE II
COMPARISON OF TOTAL MEMORY AND COMPUTATION TIME REQUIREMENTS BETWEEN A GLOBAL METRIC RBPF AND HMT-SLAM

Method	**Málaga dataset**	**Edmonton dataset**
Global RBPF	197Mb, 103min	84Mb, 39min
HMT-SLAM	36Mb, 26min	28Mb, 8min

metric mapping [26]. The performances in computation time and memory requirements are summarized in Table II. These values have been obtained for a 2.0GHz Pentium M (1Gb RAM) and for occupancy grid maps with a cell size of 12cm. It is noticeable that HMT-SLAM outperforms global RBPF for both datasets. The improvement in the memory requirements follows from the fact that in our implementation of HMT-SLAM, each particle carries, apart from the topological path hypothesis, a metric map for the current area only, while the sub-maps of previous areas are kept in a compact form [27]. In the global RBPF, each particle carries a hypotheses of the whole metric map. Therefore, the storage efficiency of a HMT map in contrast to a global RBPF becomes more and more relevant for increasingly larger environments. This reasoning is supported by the evolution of the memory requirements over time for the Málaga dataset, plotted in Fig. 9(e). Regarding the lower computation time of our approach, it is a direct consequence of the reduced number of particles (we use 15 samples). However, with only a few particles in our approach we can achieve a representation of uncertainty better than the one attainable by a global RBPF with a practical number of particles (e.g. less than 100). This turns into more precise loop closures and a more reliable representation of uncertainty.

To get an approximate idea of the time consumed by each of the processes within our implementation (described in Section V) we have measured the average time taken by each one per action-observation pair in the datasets. As can be seen in Fig. 9(c), the SLAM and AAM processes take roughly the same time in both datasets, while there are large differences for TSBI and TLCA. This reflects the fact that those processes become more time consuming for a higher number of topological areas and potential loop closures, hence they take much more time in the Málaga dataset than in the Edmonton dataset. Note that the RTL process has not been included in these experiments since they have been performed off-line. Nevertheless, it involves a negligible complexity in comparison to the rest.

It is interesting that for both datasets our method quickly converges to the correct topology after each loop closure. To visualize this process, consider the first closure of a long loop in the Málaga dataset, where the robot leaves the area labeled as ’14’ and reenters the area ’1’ (refer to Fig. 9(a)). Then, two hypotheses for the robot topological location are generated: a new, unexplored area (’15’), and an existing area that closes the loop (’1’). The intuitive idea of the expected behavior of our filter at this point is that the actual topological hypothesis will fit better to the robot observations, thus it will be assigned higher likelihood values within the particle filter that estimates the robot hybrid path s^t. Eventually, the correct

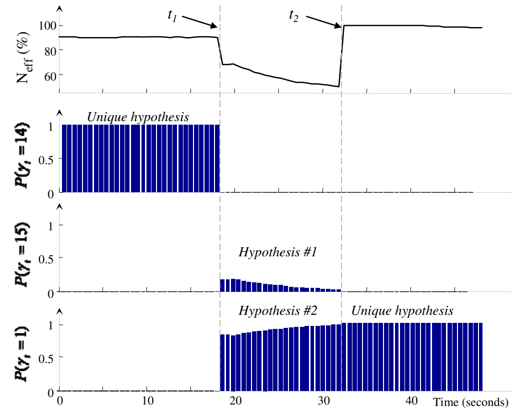

Fig. 10. The evolution over time of some values in the mapping process while closing a loop between areas ’14’ and ’1’. From the top to the bottom, the graphs represent the effective sample size (N_{eff}) of the particle filter, and the marginal PMF for the robot topological position evaluated for the areas ’14’, ’15’, and ’1’, respectively. Refer to the text for further discussion.

hypothesis will become much more likely than the rest, which will be ultimately removed by resampling. This process is clearly observable in Fig. 10, where we represent the effective sample size (ESS) [33] of the filter and the marginal PMF of the topology over time. In these graphs we can see how before time t_1 there is a unique hypothesis for the topological position (area ’14’), and next two possibilities are considered: an unexplored area ’15’ and a loop closure ’1’. Since the loop closure hypothesis provides a much better explanation of the robot observations, its probability quickly increases, and after a resampling it becomes the unique hypothesis in the filter, definitively closing the loop.

It is also desirable to contrast how other hybrid methods perform in similar loop-closure situations. The Hierarchical SLAM approach reported in [19] computes the loop closure by least-square error minimization, thus it considers just one (metric) hypothesis for the closure. Although this may be enough in many situations, in highly repetitive scenarios it would be more advantageous to maintain multiple hypotheses until the closure becomes unambiguous. This is the case of the Atlas framework [7], which considers the so-called *juvenile* hypotheses for closing loops. In that work, a juvenile hypothesis is promoted to *mature* when it performs much better than the rest, and until that point it is not allowed to modify its local map, that is, the treatment of hypotheses is purely heuristic. This is in contrast to our unified and mathematically grounded approach where there are no distinctions between all the existing topological hypotheses, and all of them are allowed to modify their corresponding local sub-maps (this, in turn, is required to perform sustainable accurate localization). Furthermore, since the Atlas framework is based on a hybrid robot pose (instead of hybrid robot *path*) it allows only one hypothesis for the robot metric pose within each sub-map, while under HMT-SLAM one can devise highly ambiguous environments where closing a long loop leads to a variety of hybrid path hypotheses, including the possibility of the robot being at the same area but with a multi-modal metric pose, and consequently with different local sub-maps.

VII. Conclusions

In this paper we have introduced a new approach for solving the problem of large-scale SLAM, which consists of the unified estimation of the hybrid metrical and topological path of the robot throughout the environment. It has been demonstrated that this idea is supported by a probabilistic structure in the SLAM problem under plausible approximations. In this paper we have addressed the probabilistic basis of a solution to HMT-SLAM in the form of sequential Bayesian filtering in the joint path-map space, which supports our approach as a promising step towards the natural integration of existing metric SLAM methods into high-level world representations, always within a Bayesian framework that manages spatial uncertainty more accurately and efficiently than previous metric and hybrid approaches. Additionally, an implementation of our ideas has been described in the form of a real-time/any-time system capable of providing an estimation of the robot pose and the map at each instant of time, giving more relevance to the computation of the metric robot pose, which may be required for navigation or manipulation purposes. Our work has been validated by experiments where relatively large HMT maps have been successfully built.

This work gives rise to a number of interesting topics that require future research, like the integration with other map types (i.e. landmarks), the simultaneous estimation of the HMT path of a team of robots, or the use of appearance and high-level knowledge for localization and loop-closure. An especially interesting issue is that of solving the robot awakening problem, or global localization, within a large-scale and *partially known* environment (a problem that can be hardly dealt with existing metric or topological methods). This issue would be faced by any mobile robot operating in a realistic lifelong application. Under the perspective of our work, the problem is naturally cast as a special case of topological loop-closure. This means that, as a byproduct, HMT-SLAM will allow a robot to incorporate to an existing map new information gathered while performing global localization, and thus unifying SLAM and global localization.

References

[1] N. Ayache and O. Faugeras, "Maintaining representations of the environment of a mobile robot," *IEEE Transactions on Robotics and Automation*, vol. 5, no. 6, pp. 804–819, 1989.

[2] T. Bailey, J. Nieto, J. Guivant, M. Stevens, and E. Nebot, "Consistency of the EKF-SLAM Algorithm," in *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2006, pp. 3562–3568.

[3] J. Blanco, "Online dataset repository," 2007. [Online]. Available: "http://babel.isa.uma.es/mrpt/downloads/"

[4] J. Blanco, J. Fernández-Madrigal, and J. Gonzalez, "A New Approach for Large-Scale Localization and Mapping: Hybrid Metric-Topological SLAM," in *Proceedings of the IEEE International Conference on Robotics and Automation*, 2007, pp. 2061–2067.

[5] J. Blanco, J. Gonzalez, and J. Fernández-Madrigal, "Consistent Observation Grouping for Generating Metric-Topological Maps that Improves Robot Localization," in *Proceedings of the IEEE International Conference on Robotics and Automation*, 2006, pp. 818–823.

[6] M. Bosse, P. Newman, J. Leonard, M. Soika, W. Feiten, and S. Teller, "An Atlas framework for scalable mapping," in *Proceedings of the IEEE International Conference on Robotics and Automation*, vol. 2, 2003, pp. 1899–1906.

[7] M. Bosse, P. Newman, J. Leonard, and S. Teller, "Simultaneous Localization and Map Building in Large-Scale Cyclic Environments Using the Atlas Framework," *The International Journal of Robotics Research*, vol. 23, no. 12, pp. 1113–1139, 2004.

[8] J. Castellanos, J. Neira, and J. Tardos, "Limits to the consistency of EKF-based SLAM," *5th IFAC Symposium on Intelligent Autonomous Vehicles*, 2004.

[9] H. Choset and K. Nagatani, "Topological simultaneous localization and mapping (SLAM): Toward Exact Localization Without Explicit Localization," *IEEE Transactions on Robotics and Automation*, vol. 17, no. 2, pp. 125–137, 2001.

[10] S. Coradeschi and A. Saffiotti, "An introduction to the anchoring problem," *Robotics and Autonomous Systems*, vol. 43, no. 2-3, pp. 85–96, 2003.

[11] M. Cummins and P. Newman, "Probabilistic Appearance Based Navigation and Loop Closing," in *Proceedings of the IEEE International Conference on Robotics and Automation*, 2007, pp. 2042–2048.

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

[13] A. Doucet, N. de Freitas, K. Murphy, and S. Russell, "Rao-Blackwellised particle filtering for dynamic Bayesian networks," in *Proceedings of the Sixteenth Conference on Uncertainty in Artificial Intelligence*, 2000, pp. 176–183.

[14] A. Doucet, S. Godsill, and C. Andrieu, "On sequential Monte Carlo sampling methods for Bayesian filtering," *Statistics and Computing*, vol. 10, no. 3, pp. 197–208, 2000.

[15] G. Dudek, M. Jenkin, E. Milios, and D. Wilkes, "Robotic exploration as graph construction," *IEEE Transactions on Robotics and Automation*, vol. 7, no. 6, pp. 859–865, 1991.

[16] H. Durrant-Whyte and T. Bailey, "Simultaneous localization and mapping: part I," *IEEE Robotics & Automation Magazine*, vol. 13, no. 2, pp. 99–110, 2006.

[17] H. Ehrig, U. Prange, and G. Taentzer, *Fundamental Theory for Typed Attributed Graph Transformation*, ser. Lecture Notes in Computer Science. Springer Berlin, 2004, vol. 3256.

[18] A. Elfes, "Using occupancy grids for mobile robot perception and navigation," *Computer*, vol. 22, no. 6, pp. 46–57, 1989.

[19] C. Estrada, J. Neira, and J. Tardos, "Hierarchical SLAM: Real-Time Accurate Mapping of Large Environments," *IEEE Transactions on Robotics*, vol. 21, no. 4, pp. 588–596, 2005.

[20] J. Fernández-Madrigal and J. Gonzalez, "Multihierarchical graph search," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 24, no. 1, pp. 103–113, 2002.

[21] U. Frese, "A Proof for the Approximate Sparsity of SLAM Information Matrices," in *Proceedings of the IEEE International Conference on Robotics and Automation*, 2005, pp. 331–337.

[22] ——, "A Discussion of Simultaneous Localization and Mapping," *Autonomous Robots*, vol. 20, no. 1, pp. 25–42, 2006.

[23] C. Galindo, J. Fernández, and J. Gonzalez, "Hierarchical Task Planning through World Abstraction," *IEEE Transactions on Robotics*, vol. 20, no. 4, pp. 667–690, 2004.

[24] C. Galindo, A. Saffiotti, S. Coradeschi, P. Buschka, J. Fernández-Madrigal, and J. Gonzalez, "Multi-hierarchical semantic maps for mobile robotics," in *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2005, pp. 2278–2283.

[25] J. González, A. Ollero, and A. Reina, "Map Building for a Mobile Robot equipped with a Laser Range Scanner," in *Proceedings of the IEEE International Conference on Robotics and Automation*, 1994, pp. 1904–1909.

[26] G. Grisetti, C. Stachniss, and W. Burgard, "Improved Techniques for Grid Mapping With Rao-Blackwellized Particle Filters," *IEEE Transactions on Robotics*, vol. 23, pp. 34–46, Feb 2007.

[27] G. Grisetti, G. Tipaldi, C. Stachniss, W. Burgard, and D. Nardi, "Fast and accurate slam with rao-blackwellized particle filters," *Robotics and Autonomous Systems*, vol. 55, no. 1, pp. 30–38, 2007.

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

[29] A. Howard and N. Roy, "The robotics data set repository (radish)," 2003. [Online]. Available: http://radish.sourceforge.net/

[30] K. Kouzoubov and D. Austin, "Hybrid topological/metric approach to SLAM," in *Proceedings of the IEEE International Conference on Robotics and Automation*, vol. 1, 2004.

[31] B. Kuipers and Y. Byun, *A Robot Exploration and Mapping Strategy Based on a Semantic Hierarchy of Spatial Representations*. Artificial Intelligence Laboratory, the University of Texas at Austin, 1990.

[32] B. Kuipers, J. Modayil, P. Beeson, M. MacMahon, and F. Savelli, "Local metrical and global topological maps in the hybrid spatial semantic hierarchy," in *Proceedings of the IEEE International Conference on Robotics and Automation*, vol. 5, 2004, pp. 4845–4851.

[33] J. Liu, "Metropolized independent sampling with comparisons to rejection sampling and importance sampling," *Statistics and Computing*, vol. 6, no. 2, pp. 113–119, 1996.

[34] Y. Liu and S. Thrun, "Results for outdoor-SLAM using sparse extended information filters," in *Proceedings of the IEEE International Conference on Robotics and Automation*, vol. 1, 2003, pp. 1227–1233.

[35] F. Lu and E. Milios, "Globally Consistent Range Scan Alignment for Environment Mapping," *Autonomous Robots*, vol. 4, no. 4, pp. 333–349, 1997.

[36] R. Martinez-Cantin, N. de Freitas, and J. Castellanos, "Analysis of Particle Methods for Simultaneous Robot Localization and Mapping and a New Algorithm: Marginal-SLAM," in *Proceedings of the IEEE International Conference on Robotics and Automation*, 2007, pp. 2415–2420.

[37] O. Martínez-Mozos, R. Triebel, P. Jensfelt, A. Rottmann, and W. Burgard, "Supervised semantic labeling of places using information extracted from sensor data," *Robotics and Autonomous Systems*, vol. 55, no. 5, pp. 391–402, 2007.

[38] J. Modayil, P. Beeson, and B. Kuipers, "Using the topological skeleton for scalable global metrical map-building," in *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems*, vol. 2, 2004, pp. 1530–1536.

[39] M. Montemerlo, S. Thrun, D. Koller, and B. Wegbreit, "FastSLAM: A factored solution to the simultaneous localization and mapping problem," *Proceedings of the AAAI National Conference on Artificial Intelligence*, pp. 593–598, 2002.

[40] ——, "FastSLAM 2.0: An improved particle filtering algorithm for simultaneous localization and mapping that provably converges," in *Proceedings of the International Joint Conference on Artificial Intelligence*, 2003.

[41] H. Moravec and A. Elfes, "High resolution maps from wide angle sonar," in *Proceedings of the IEEE International Conference on Robotics and Automation*, vol. 2, 1985.

[42] H. Moravec, "Sensor Fusion in Certainty Grids for Mobile Robots," *AI Magazine*, vol. 9, no. 2, pp. 61–74, 1988.

[43] K. Murphy, "Bayesian map learning in dynamic environments," *Advances in Neural Information Processing Systems (NIPS)*, vol. 12, pp. 1015–1021, 1999.

[44] J. Neira and J. Tardós, "Data association in stochastic mapping using the jointcompatibility test," *IEEE Transactions on Robotics and Automation*, vol. 17, no. 6, pp. 890–897, 2001.

[45] L. Paz, P. Jensfelt, J. Tardós, and J. Neira, "EKF SLAM Updates in O(n) with Divide and Conquer SLAM," in *Proceedings of the IEEE International Conference on Robotics and Automation*, 2007, pp. 1657–1663.

[46] A. Ranganathan, E. Menegatti, and F. Dellaert, "Bayesian Inference in the Space of Topological Maps," *IEEE Transactions on Robotics*, vol. 22, no. 1, pp. 92–107, 2006.

[47] E. Remolina and B. Kuipers, "Towards a general theory of topological maps," *Artificial Intelligence*, vol. 152, no. 1, pp. 47–104, 2004.

[48] D. Rubin, "Using the SIR algorithm to simulate posterior distributions," *Bayesian Statistics*, vol. 3, pp. 395–402, 1988.

[49] S. Russell and P. Norvig, *Artificial intelligence: a modern approach, chapter 14*. Prentice-Hall, Inc. Upper Saddle River, NJ, USA, 1995.

[50] F. Savelli and B. Kuipers, "Loop-Closing and Planarity in Topological Map-Building," in *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems*, vol. 2, 2004, pp. 1511–1517.

[51] J. Shi and J. Malik, "Normalized Cuts and Image Segmentation," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 22, no. 8, pp. 888–905, 2000.

[52] R. Smith, M. Self, and P. Cheeseman, "A stochastic map for uncertain spatial relationships," *The fourth international symposium on Robotics Research*, pp. 467–474, 1988.

[53] ——, "Estimating uncertain spatial relationships in robotics," *Autonomous Robot Vehicles*, vol. 1, pp. 167–193, 1990.

[54] C. Stachniss, G. Grisetti, and W. Burgard, "Recovering Particle Diversity in a Rao-Blackwellized Particle Filter for SLAM After Actively Closing Loops," in *Proceedings of the IEEE International Conference on Robotics and Automation*, 2005.

[55] C. Stachniss, D. Hahnel, and W. Burgard, "Exploration with active loop-closing for FastSLAM," in *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems*, vol. 2, 2004.

[56] J. Tardos, J. Neira, P. Newman, and J. Leonard, "Robust Mapping and Localization in Indoor Environments Using Sonar Data," *The International Journal of Robotics Research*, vol. 21, no. 4, pp. 311–330, 2002.

[57] S. Thrun, *Robotic Mapping: A Survey*. School of Computer Science, Carnegie Mellon University, 2002.

[58] ——, "Learning Occupancy Grid Maps with Forward Sensor Models," *Autonomous Robots*, vol. 15, no. 2, pp. 111–127, 2003.

[59] S. Thrun and A. Bucken, "Integrating grid-based and topological maps for mobile robot navigation," *Proceedings of the Thirteenth National Conference on Artificial Intelligence (AAAI-96)*, pp. 944–950, 1996.

[60] S. Thrun, W. Burgard, and D. Fox, *Probabilistic Robotics*. The MIT Press, September 2005.

[61] N. Tomatis, I. Nourbakhsh, and R. Siegwart, "Hybrid simultaneous localization and map building: a natural integration of topological and metric," *Robotics and Autonomous Systems*, vol. 44, no. 1, pp. 3–14, 2003.

[62] T. Verma and J. Pearl, "Causal networks: Semantics and expressiveness," in *Uncertainty in Artificial Intelligence 4*, R. D. Shachter, T. S. Levitt, L. N. Kanal, and J. F. Lemmer, Eds. Amsterdam: North-Holland, 1990, pp. 69–76.

[63] Z. Zivkovic, B. Bakker, and B. Krose, "Hierarchical map building using visual landmarks and geometric constraints," in *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2005, pp. 2480–2485.

Jose-Luis Blanco was born in Linares (Spain) in 1981. He received the M.S. degree in Electrical Engineering from the University of Málaga in 2005. He is currently a PhD student in the Department of System Engineering and Automation at that University. He is the author of about 20 journal and conference papers. His research interest includes robot autonomous navigation, world modeling, and computer vision.

Juan-Antonio Fernández-Madrigal was born in Córdoba (Spain) in 1970. He received the M.S. degree in computer science in 1994, and the PhD in 2000 at the University of Málaga (Spain), where he is currently associate professor (tenured). He is author of near 40 journal and conference papers and two books on multiple abstraction applied to mobile robots. His main research interests during the last ten years have been the use of abstraction for real robots and robotic software development frameworks.

Javier González received the B.S. degree in Electrical Engineering from the University of Seville in 1987. He joined the Department of System Engineering and Automation at the University of Malaga in 1988 and received the Ph.D. from this University in 1993. In 1990-1991 he was at the Field Robotics Center, Robotics Institute, Carnegie Mellon University (USA) working on mobile robots. Currently, he is an associate professor at the University of Malaga and he has led several Spanish and European projects on mobile robot from 1996. His research interest includes mobile robot autonomous navigation, computer vision and remote sensing. In these areas, he is author and co-author of more than 80 papers and three books about computer vision and mobile robot symbolic world modelling.