IROS 2004. http://www.cs.utexas.edu/users/qr/papers/Savelli-iros-04.html

Loop-Closing and Planarity in Topological Map-Building

Francesco Savelli
Dipartimento di Informatica e Sistemistica
Università di Roma "La Sapienza"
Roma, Italy
Email: savelli@dis.uniroma1.it

Benjamin Kuipers
Department of Computer Sciences
University of Texas at Austin
Austin, Texas USA
Email: kuipers@cs.utexas.edu

***Abstract*— Loop-closing has long been recognized as a critical issue when building maps of large-scale environments from local observations. Topological mapping methods abstract the problem of determining the topological structure of the environment (i.e., how loops are closed) from the problem of determining the metrical layout of places in the map and dealing with noisy sensors. A recently developed incremental topological mapping algorithm [1], [2] generates all possible topological maps consistent with the experienced sequence of actions and observations and the topological axioms. These are then ordered by a preference criterion such as minimality or probability, to determine the single best map for continued planning and exploration. This paper presents the planarity constraint and analyzes its impact on the search-tree of all topological maps consistent with (non-metrical) exploration experience. Experimental studies demonstrate excellent results even in artificial environments where loop-closing is particularly difficult due to large amounts of perceptual aliasing and structural symmetry.**

I. Introduction

When a robot or human agent explores an unknown environment and builds a map, a key type of decision is whether the current place is the same as a previously-mapped place or a new one. This decision determines which loops exist in the map.

The loop-closing problem is a particularly difficult case of the data association problem for simultaneous localization and mapping (SLAM) algorithms building metrical maps in single global frames of reference [3]. The number and identity of the visited places and observed landmarks determine part of the conditional dependency relations in the probabilistic model — usually a dynamic Bayes net — wherein an instance of the SLAM problem is traditionally solved when data association is known *a priori*. This type of uncertainty can thus be addressed as the problem of determining the right structure of the Bayes net representing the SLAM problem at hand; in this respect, uncertainty on the Bayes net's topology reflects topological ambiguity in the physical world (data association).

Loop-closing has also long been recognized as an important problem in the topological mapping literature [4], [5], [6].

In this paper, we present and analyze the *planarity constraint*: the requirement that the topological map be embedded in the plane without crossing edges. Our focus is on the topological mapping process, which takes a symbolic description of the environment exploration as input. How such a symbolic description can be extracted from raw sensor-data is beyond the scope of this paper. An approach based on the use of local metrical maps for real robots equipped with laser range-finder sensors is discussed elsewhere [2], along with experiments in indoor environments. Similar and alternative techniques for place detection and topological structure extraction in occupancy grid-maps were also proposed in [7] and [8], [9].

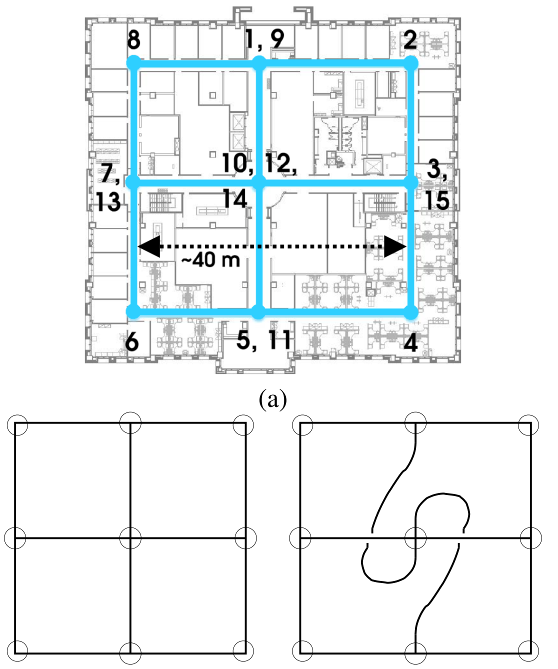

(a)

(b) (c)

Fig. 1. **Unintended models.** Both (b) and (c) are among the consistent topological hypotheses after the complete physical robot exploration of the environment (a). Places and their local perceptual characteristics at the topological abstraction level are represented by the small circles. Due to perceptual aliasing of places and the symmetrical structure of the environment, any further exploration will give the same sequence of observations, making it impossible to discriminate between the two maps. Planarity constraint is an inexpensive alternative to using metrical information or raw-data matching to discard (c), and applies directly to the topological abstraction. Indeed (b) and (c) represent the same graph, but considered as *embedded* graphs, (b) is planar, while (c) is not.

Figure 1 illustrates a physical exploration scenario addressed in our previous work, which motivates the application of the planarity constraint. In this case, the unintended consistent (but non-planar) model could be also discarded using odometrical information to enforce global metrical consistency by means of state-of-the-art SLAM methods. However, the planarity test has the advantage of being independent of the actual geometric properties and metrical scale of the real environment, as it yields a boolean result on the available abstract topological representation. Moreover, it can be performed in time and space *linear* in the number of topological links. If it is dynamically evaluated as the map grows during the topological search —exploiting past computations when a new element is inserted— a *sublinear* ($O(log^2(n))$ worst-case time) algorithm can be used [10]. We show that the planarity constraint significantly constrains the search for correct topological maps in particularly difficult types of environment, where the agent is faced with poor perceptual characterization of distinctive locations and landmarks, large amount of perceptual aliasing, multiple nested loops, and structural symmetry. In some relevant cases, the planarity test makes it possible to find unique optimal maps.

II. Topological Mapping

The Spatial Semantic Hierarchy (SSH) [11] describes large-scale space using four different levels of representation. At the *control* level, the agent repeatedly selects a hill-climbing control law to converge to and localize at a *distinctive state* (dstate), and then a trajectory-following control law to move from the current dstate to the neighborhood of another, where hill-climbing converges to the next dstate, eliminating cumulative error. The *causal* level abstracts this pattern of behavior to a deterministic automaton, consisting of states (the distinctive ones), *actions* (sequences of control laws), *schemas* $\langle x, a, x' \rangle$ (asserting that state x' results from performing action a in state x), and *views* (the perceptual images of states, $view(x, v)$). The *topological* level distinguishes between turn and travel actions, and aggregates states into places, paths, and regions, related by connectivity, order, and containment. The *metrical* level consists of local metrical attributes annotating objects at the causal and topological levels, local metrical models of small-scale space in place neighborhoods, and (when resources permit) global metrical models of the large-scale environment [12]. A formalization of the topological map in non-monotonic logic, and an algorithm for identifying minimal models according to a prioritized circumscription policy is given in [1].

Hybrid mapping methods based on the SSH allow metrical maps of local regions to be linked into topological maps of large-scale space. Increasingly efficient algorithms have been developed to exploit structure obtained from local metrical models [1], [13], [2]. These algorithms generate all possible models (i.e., topological maps), filter out those inconsistent with the topological axioms, and provide a preference ordering on the remaining models.

Exploration experience is an alternating sequence of actions and views. Repeat the following for each action a and resulting view v from the beginning to the end of the sequence.
For each $\langle M, x \rangle$ on the fringe of the tree:

1) If M includes a schema $\langle x, a, x' \rangle$
   then let v' be such that $view(x', v')$,
   - if $v' = v$, then $\langle M, x' \rangle$ is the successor to $\langle M, x \rangle$;
   - if $v' \neq v$, then mark $\langle M, x \rangle$ as inconsistent.

   (All the views corresponding to successive turn actions in the same local place neighborhood are known all at once from a local metrical map, and include all the gateways for leaving/approaching the place. Thus if a is a turn action, M always contains $\langle x, a, x' \rangle$.)
2) Otherwise, M does not include $\langle x, a, x' \rangle$. This means that x is a "pending gateway" of its place (see next section), a is a travel action, and x' must be a pending gateway of the place it is associated with.
   Let M' be M extended with a new distinctive state symbol x' and the assertions $view(x', v)$ and $\langle x, a, x' \rangle$. Consider the $k \geq 0$ dstates x_j in M such that:
   a) $view(x_j, v)$, and
   b) x_j is a pending gateway of its place, and
   c) x_j must belong to the same face of the graph as x (this accounts for the planarity constraint explained in the next section).

   Then $\langle M, x \rangle$ has $k + 1$ successors:
   - $\langle M' \cup \{x' = x_j\}, x' \rangle$ for $1 \leq j \leq k$, plus
   - $\langle M' \cup \{\forall j \; x' \neq x_j\}, x' \rangle$.
3) If any of the new successor maps violates the topological axioms, mark it inconsistent.

After each action a and resulting view v, the nodes $\langle M, x \rangle$ at the leaves of the search-tree are the maps and qualitative poses that are consistent with exploration experience thus far.

Fig. 2. **Building the tree of topological maps.**

We briefly introduce an algorithm (Figure 2) that constructs a tree of topological hypotheses. Its input is the abstract (non-metrical) description of actions and views summarizing the exploration experience at the topological level. Tree nodes are pairs $\langle M, x \rangle$, where M is a map and x is the dstate representing the qualitative pose (vertex-position and edge-orientation of the underlying graph) within the topological map. After each action, the pair $\langle M, x \rangle$ is linked to its successor(s) $\langle M', x' \rangle$. If the action takes place entirely within the existing map, then $M' = M$. If otherwise x' is a new dstate or if the link from x to x' is new, then M' is an extension of M and it must be checked for consistency with the topological axioms and possibly identified as inconsistent. Most importantly for loop-closing, if x' has the same view as some existing states, it might be the same as one of those states or it

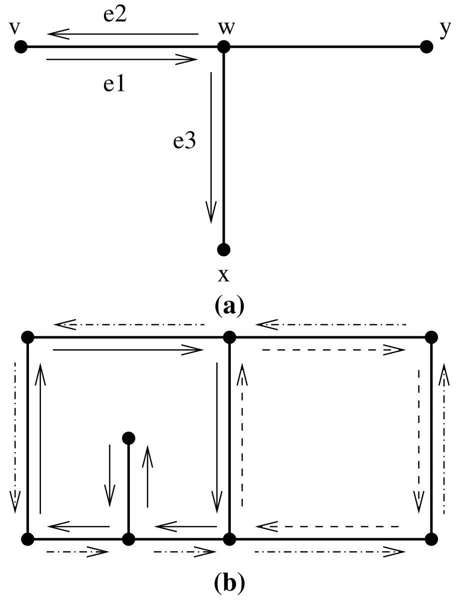

Fig. 3. **Face Tracking** If edges are imagined as corridors, face tracking can be intuitively viewed as walking along loops always following the right wall. **(a)** $e2 = rev(e1)$, $e3 = pred(e2) = pred(rev(1)) = next(e1)$. **(b)** the three faces of a simple embedded graph.

might be a new state, so the tree of maps must branch on all possibilities.

Thus, the leaves of the map-tree represent the possible consistent maps and poses at the current time. A preference ordering, currently based on minimality as defined by the circumscription policy, defines the optimal map(s) on which to base planning for the next steps in exploration or navigation. The preference ordering can be a partial order, so if there are multiple distinct optimal maps, further exploration is necessary to discriminate among them. The algorithm can be strengthened either with stronger constraints to rule out more maps as inconsistent, or with a stronger preference ordering.

Unfortunately, as shown in Figure 1, there are unintended models (incorrect maps) which are consistent with the topological axioms as well as all past and all possible future experiences during exploration. The planarity constraint (2c in Figure 2) significantly helps to rule them out.

III. The Planarity Constraint

A. Embedded graphs and planarity

We recall here basic mathematical concepts of graph theory [14].

An *undirected graph* is intended, as usual, as a set of vertices and undirected edges $G = (V, E)$. An *embedding* of an undirected graph G defines a clockwise circular order over all the edges incidents on a vertex v, around each v of G. An *embedded graph* is an undirected graph equipped with an embedding. When embedded graphs are studied as special structures themselves, they are also called *maps*. Although this terminology witnesses how such structures are mathematically relevant to our purposes, we use the term 'map' in the more general sense intended in the SSH.

The *genus* of an embedded graph G is the minimal genus of an orientable surface on which G can be drawn enforcing its embedding without edge-crossings. In this case G is *planar* on the considered surface. When G is planar on the plane, its genus is 0 and we simply say it is planar. (Note that the plane and the sphere both have genus 0, so G is planar iff it is so on the sphere too. The torus is a surface with genus 1.)

We describe here the classical method for deciding the planarity of an embedded graph.

For any embedded graph $G = (V, E)$, consider the transformation T such that:

- $T(G) = (V, T(E))$ is a bidirected graph obtained replacing each undirected edge $e \in E$ that links vertices v and w by two directed edges e_v and e_w where $source(e_v) = target(e_w) = v$ and $source(e_w) = target(e_v) = w$. We define the $1-1$ mapping $rev(e_v) = e_w$, $rev(e_w) = e_v$ (note that $rev(rev(x)) = x$).
- The circular order of undirected edges $[e^1 \dots e^k]$ around each vertex v defined by the embedding in G is reflected by the circular order $[e_v^1 \dots e_v^k]$ of outward directed edges leaving v in $T(G)$.

The mapping $succ : T(E) \rightarrow T(E)$ associates to any directed edge e_v the next outward directed edge in the circular clockwise order around v. Analogously for $pred : T(E) \rightarrow T(E)$, counterclockwise.

Computing the genus of an embedded graph G requires us to count its *faces*, which we define on the auxiliary structure $T(G)$. Consider the following function $next : T(E) \rightarrow T(E)$ such that $next(e_v) = pred(rev(e_v))$, i.e., the directed edge $next(e_v)$ is the predecessor of $rev(e_v) = e_w$ in the circular order of edges around $w = target(e_v)$ (Figure 3 (a)). Note that $next(e_v)$ always exists and is unique for any $e_v \in T(E)$. Now let $[e_1, e_2 \dots e_n]$ be a list of directed edges of $T(G)$ s.t. (i) $e_{i+1} = next(e_i)$ for all $i = 1 \dots n-1$, (ii) $e_i \neq e_j$ for any $1 \leq i < j \leq n$, and (iii) $next(e_n) = e_1$. We call the directed loop formed by the sequence of edges in such a list a *face* (Figure 3 (b)).

The preceding paragraph provides a combinatorial definition of faces. If we consider a surface on which an embedded graph is planar, we can geometrically characterize faces as the partitions induced on the surface by the images of the edges. The infinite face of a planar embedded graph on the plane can be avoided if the drawing is considered on the sphere. Faces can then also be intuitively viewed as the inside of the "least" loops, that is, those loops that do not contain other loops.

The genus of an embedded graph G is given by the following (Euler-Poincaré) formula:

$$genus(G) = (|E| + 2c - |V| - |I| - f)/2,$$

where I is the set of isolated vertices, c the number of connected components, and f the number of faces. Since we shall only deal with totally connected ($c = 1$, $|I| = 0$) embedded graphs — they model maps from single-agent explorations — planarity (genus=0) can be tested [1] by:

$$|V| - |E| + f = 2,$$

the well known Euler's formula on the number of faces in a polyhedron.

The number of faces can be counted in linear time and space (in the number of edges) following the mathematical formulation given above, based on the function $next$.

B. Search space reduction

The number of maps that the planarity constraint discards among all the topological hypotheses depends on many factors, including the nature and structure of the environment, the number and location of perceptually aliased places and the particular exploration route. Nevertheless, it is possible to formalize the reduction of the branching factor at a node in the search-tree, in a fashion that sheds light on the role of the past choices that have shaped the map at that node.

Each place has a number of *gateways*. An edge is created by linking gateways at two places. In general, while exploration is underway, some places have *pending gateways* that have not yet been explored and linked to make edges. Since pending gateways are not yet edges of the embedded graph, they must be skipped by the functions $succ$, $pred$, and $next$ while counting the faces (even though they predict the positions in the circular order around their vertices of the edges that will eventually connect them). A map is *closed* if it has no pending gateways. A closed map implicitly represents the entire environment.

Every pending gateway in a non-closed map belongs to exactly one face of the underlying embedded graph. We formalize this notion as follows. Consider any adjacent pair $e_{i+1} = next(e_i)$ of directed edges of the loop defining a face, and let v be the vertex $v = target(e_i) = source(e_{i+1})$. Any pending gateway pe of v that comes after e_{i+1} and before $rev(e_i)$ in the circular clockwise order of outward directed edges around the vertex v *belongs to the face*. (When only one edge is currently incident on v, i.e., when $e_{i+1} = rev(e_i)$, then any pending gateway of v trivially belongs to the face.) It follows that the set P of all the pending gateways in a map is partitioned into the sets $P_1 \dots P_f$, each P_i collecting the pending gateways belonging to face i, $i = 1 \dots f$, where f is the number of faces. Let $p_i = |P_i|$ and $p = |P|$.

Now, consider a node of the search tree where the next exploration step leaves the current place through one of its pending gateways. The branching factor of this node is given by the number of possible successors of this map, consistent with the resulting observation. One successor is always the map where the arrival is at a totally new place. The others are maps where the move is to a place that is already represented in the map. First, such a place must have been approached through one of its pending gateways, and the resulting view must be consistent with the observation made. Second, if planarity is assumed, *the two pending gateways to be linked must belong to the same face*, otherwise they would be unified in an edge that would cross the loop defining the face.

Assuming recursively that the map at the current node of the search tree is already planar, we are interested in the reduction of the branching factor when planarity is enforced also in its children. In particular we consider the ratio b_p/b where b_p is the number of planar children of this node and b the number of all the children, including the non-planar ones.

Assume a worst-case scenario where every pending gateway in the map is compatible with the next observation, and is thus a potential arrival. Then b is equal to p, minus the one pending gateway where the move starts, plus the map with a totally new place, i.e., $b = p$. Analogously, if the starting pending gateway belongs to the ith face, $b_p = p_i$. Then $b_p/b = p_i/p$. Denoting the average number of pending gateways per face $\bar{p} = p/f$, we have $b_p/b = p_i/(\bar{p} \cdot f)$.

From the recursive assumption of planarity, we have $f = |E| - |V| + 2$ (before including the new edge in E). The interesting point here is that the difference $|E| - |V|$ depends on the choices made at past branching points leading to this map. Every time it was decided to explain the observation at hand by linking the place with one already encountered, unifying two pending gateways, $|E| - |V|$ increased by one unit. If instead it was decided to build a totally new place, then a vertex and an edge were inserted at once, and $|E| - |V|$ did not change. Considering that at the root of the search tree there is a map with just one vertex and no edge, $|E| - |V| = m - 1$ where m is the number of matches made between existing places, as in the first case above. Thus, finally, the dynamic ratio that quantifies the branching factor reduction is[2]:

$$\frac{b_p}{b} = \frac{p_i}{\bar{p}} \cdot \frac{1}{m+1}, \tag{1}$$

The first factor normalizes according to the relative number of consistent pending gateways in the current face w.r.t. the average face. The second factor confirms the intuition that the more loops have been closed, the more topologically compact the map must be, and therefore the fewer ways there are to close new loops while preserving planarity.

IV. Experimental Results

We have carried out several experiments to investigate in detail how planarity testing improves the topological search.

[1] If unlikely non-planar maps must be taken into account, the planarity constraint might be relaxed, by simply preferring maps with smaller genus.

[2] Note that we get the same result if we replace the worst-case assumption that all pending gateways are consistent with the observation with a more moderate assumption that pending gateways consistent with the observation are uniformly distributed throughout the map. In this case, we restrict the numbers p_i, p, and $\bar{p}$ to only those pending gateways consistent with the observation.

Environment Structure:	3×3			3×4			4×4 (up to 16 places)		
Planarity Pruning:	no	yes	$\sim$ red.	no	yes	$\sim$ red.	no	yes	$\sim$ red.
Search-tree nodes built:	317	240	24%	4,239	1,722	59%	192,569	8,452	95%
Maps built:	188	147	22%	3,400	1,425	58%	93,034	4,951	94%
Final maps (all):	88	62	29%	1,465	507	65%	23,685	545	97%
Final maps (closed):	6	3	50%	16	3	81%	6992	58	99%
Optimal Final Maps:	2	1	/	8	1	/	2	2	/

Fig. 4. **Search Space Reduction** Breadth-first search expands the entire space of topological hypotheses permitting to observe the exact reduction of the search-space if planarity is assumed. The 4×4 results regard only the portion of space of maps with less than 17 places, and show that the planarity constraint has greater impact as more loop-closures are forced. Since minimality is related to loop-closing, the planarity constraint will prove especially useful in practical cases when a best-first search is employed.

The topological map-builder does not assume a particular global structure for environment. However, for any topological mapping algorithm, the worst-case environments will be those with large amounts of perceptual aliasing (different places that look the same) and structural symmetry (because it may be difficult or impossible to refute incorrect hypotheses).

We evaluate this algorithm using simulated square and rectangular grids, to maximize the difficulty facing the algorithm. The only places are corridor intersections, with "L", "T", or "+" structure. Every pose is perceptually aliased multiple times throughout the environment and the global structure of the environment has multiple embedded loops and strong symmetries along multiple axes. This kind of abstract environment allows a fair and straight comparison of experimental results as the environment scales topologically. Besides, its pattern is relevant to several real environments, such as outdoor urban layouts, or indoor large libraries with long corridors and shelves that strongly limit the agent's sensory horizon in most of the locations.

The results we report are for a "snake" pattern of exploration. The agent starts from a corner and walks along all the horizontal corridors, in alternate directions, moving from a corridor to the next parallel one when it reaches a "T" (just a corner the first time) intersection. Then it starts an analogous exploration of the vertical corridors. We believe that this kind of exploration would also prove hard when a pure metrical mapping method is used that closes loops based on a maximum-likelihood choice. Indeed, if corridors are long enough, the inevitable angular odometrical error might often support the hypothesis that the agent is back to a "+" intersection on a parallel corridor. In this case, if maximum-likelihood is used for a greedy on-line search, the map would be irreversibly affected by such an error.

We assume the agent can acquire the correct abstract representation "L", "T", or "+" of the place being visited and be able to discriminate between these views.

We have run the experiments on an implementation of the SSH topological breadth-first search, with and without planarity testing. The reason for experimenting with breadth-first search is that it builds the complete set of the current topological hypotheses on-line: this allows us to observe the exact reduction of the whole search-space when planarity pruning is applied. We have gathered statistics about (1) the number of search-tree nodes, (2) maps ever built (these are usually fewer than search-tree nodes because when a map already correctly predicts the result of an action, only the assumed current position changes from node to node), (3) the number of final maps (those on the consistent leaves of the search-tree at the end of the exploration), (4) final maps that are also closed, and (5) final maps that are also optimal (according to the preference defined by the circumscription). The table in Figure 4 collects the results for the environments discussed below.

The first column illustrates the results for a 3×3 grid (see Figure 1(a)). Planarity testing discards half of the final closed maps, and allows to determine the correct model as the unique optimal final map. (The two optimal maps obtained when planarity is not assumed are those in Figure 1.)

The second column presents results from a single complete exploration of a more topologically complex environment: a 3×4 grid of places. The gain in reducing the search is larger than in the single exploration of the 3×3 grid. This is due to the branching-factor reduction, which makes planarity testing *exponentially more efficient* as the exploration proceeds.

In the last column we address a 4×4 grid. Note that, since now the two axes have same length, the environment has one more degree of symmetry, and the chances of ruling out wrong models by inconsistencies along observation-sequences is much reduced. A breadth-first search of the maps consistent with a complete exploration does not terminate in reasonable time. However, we have driven the search so as to gather statistics significant for our purposes. We have asserted an upper bound of 16 on the number of places a map can have. In our case this was intentionally chosen equal to the actual number of places in the assumed physical environment. During the search any partial map that grows over this limit is discarded. Since during the search the number of places in a partial map never decreases, *no final map with less than 17 places that would be produced with the complete search is lost*. (Note

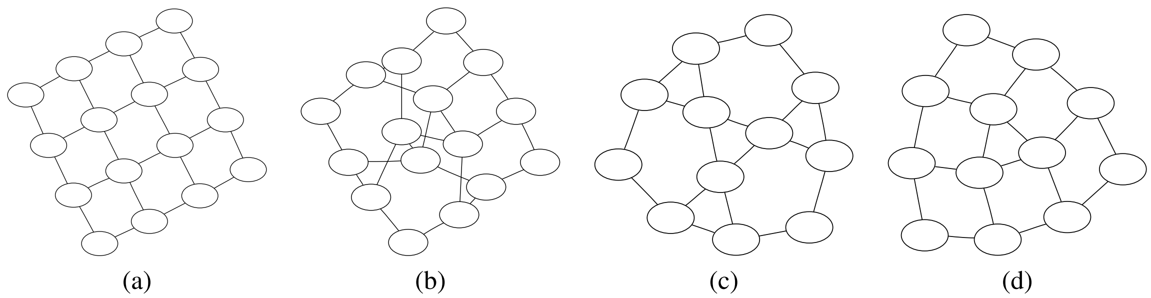

Fig. 5. 4×4 **grid models** After the "snake" exploration of the 4×4 grid, and considering only closed and qualitatively perpendicular maps, (a) is the unique (and correct) planar map. (b) is one example of the 255 non-planar "relatives" of (a) that are filtered out by the planarity constraint. Without the qualitative perpendicular requirement, there are two consistent "over-minimal" maps with only 12 states, (c) and (d). Sufficient further exploration can rule out (c), but not (d).

that having an upper bound — not necessarily picking the exact correct number of places — could be a reasonable strategy in some practical situations. The completeness of the search under this limit would then prove useful.)

Besides the previous considerations we can observe here that as the search proceeds, only those maps that account for a certain minimum number of loop-closures can remain under the upper-bound and not be ruled out. This means that the number of matches m in the denominator of the dynamic ratio formalized in Equation (1) must grow at a certain average rate. Therefore, the dramatic performance of planarity pruning in this last experiment is coherent with our formalization, and with the intuition that in more compact maps there are far fewer ways to allow spatially for a new link while avoiding edge-crosses.

The two optimal final maps found with planarity pruning are "overminimal". That is, due to symmetry and aliasing in the actual environment, the maps are complete and consistent with exploration experience, but they have fewer places than the actual environment (Figure 5(c,d))[3].

We went further in making sense of the 4×4 grid exploration data. Among *all the final maps* we have considered those closed maps that are "qualitatively perpendicular", i.e., no two global paths intersect each other in more than one place and no three global paths form a triangle. There is only one such map (the correct model of the environment) in the case of planarity pruning, as opposed to 256 such maps in the other case (all with 16 places). The unique planar solution is shown in Figure 5(a), while a map among the 255 non-planar ones is shown in Figure 5 (b). These numbers provide a concrete insight into how many non-planar "close relatives" of the correct topological model can arise in such a symmetrical environment.

We have observed that a *best-first* search, where the search-tree expansion is prioritized by the optimality policy, in some cases produces the same results while reducing the time and space used dramatically. Furthermore, since minimality of the number of places is a component of the optimality policy, by similar reasoning as with the upper bound above, we expect the maps thus prioritized to be those where m is larger and so the planarity constraint tends to do more work (by Equation (1)). That is, best-first search better leverages the potential of the planarity assumption. Applied to the snake-exploration of the 4×4 grid *without upper bound on the number of places*, with pruning of non-planar and non-qualitatively-perpendicular maps, and with backtracking if the final solution is not closed, our C++ implementation of best-first search determines the correct model in 0.84 sec. on an Intel Pentium 1.5 Ghz. However, it does not always ensure similar benefits over breadth-first search as the exploration/environment grows in size and topological ambiguity.

The relationship between the particular preference policy chosen and optimal-completeness of best-first search (which is guaranteed by exhaustive searches such as breadth-first) requires more investigation (although a consistent solution is always found, if one exists).

V. Related Work

Following [15] where the topological nature of cognitive maps is pointed out, embedded graphs as representations of topological maps were proposed in [5]. They show that correct map-building is impossible in general if only the cyclic order of the incident edges is used to recognize a place, unless the agent is provided with a portable marker it can drop and pick up. In this case the agent can cope with perceptual aliasing and symmetries, and learn the correct topological model. The upper bound on the length of the exploration with a portable marker reduces from polynomial to almost linear in the size of the graph, when planarity is assumed [16]; a similar improvement is obtained also for the map-validation problem [17].

The approach presented in [18] is closer to our work in that there is no recourse to portable markers. It provides an algorithm that expands a tree of hypotheses about loop-closures, but the role of planarity in branching-factor reduction is not investigated, and a rich spatial ontology that allows for a preference policy is not used.

Because of the negative result above, several works have addressed topological map-building as the problem of learning the minimal "discernable" structure of the environment, i.e., its smallest underlying automaton: for

[3]Note that this example illustrates the kind of extreme structural symmetry that requires a portable marker to find the correct map, as in [5].

example, in [19] the deterministic case is addressed, while the case of stochastic/noisy observations is introduced in [20]. Here, the minimal automaton representation and other assumptions are not compatible with environments such as grids. Despite the "general impossibility" of dealing with such ambiguous topologies, our work shows that it is possible to handle them to some interesting extent.

VI. Conclusions and Future Work

The planarity constraint allows to filter out most incorrect topological loop-closure hypotheses inexpensively, and independently of the actual metrical scale and geometrical appearance of the loops at issue. Although we have focused on topological mapping in isolation, our study suggests that this constraint would prove very useful also in hybrid metrical-topological mapping methods [21], [22].

The contribution of our work is twofold. First, concerning a general on-line mapping algorithm that does not assume a particular exploration strategy [18], [2], we have provided formal and experimental analysis of the great impact of planarity constraint on the search-space. Second, in so doing, we have shown how topological map-building can go considerable distance in dealing with embedded loops, perceptual aliasing, and symmetry, when exploration experience is interpreted in terms of a rich spatial ontology [11] that allows to rule out inconsistent models, or to define a preference policy.

The cases we have considered here are extreme in the weak perceptual characterization and adversity of the environments being explored. They are meant to investigate the potential of topological mapping. In real cases, *any* kind of additional feature or irregularity in the structure can be exploited to test and refute those incorrect models that we have assumed impossible to discriminate (see for example [2], [23]).

Finally, we plan to combine topological mapping with modern SLAM metrical algorithms. The contribution from topological mapping as presented in this paper would be the strong reduction of data association uncertainty — one of the most critical issue in SLAM — while odometrical information could be used to improve the selection of the final map, or to better prioritize nodes expansion in the search-tree, by a best-first strategy.

Acknowledgment

This work has taken place in part in the Intelligent Robotics Lab at the Artificial Intelligence Laboratory, The University of Texas at Austin. Research of the Intelligent Robotics lab is supported in part by the Texas Higher Education Coordinating Board, Advanced Technology Program (grant 003658-0656-2001), and by an IBM Faculty Research Award. Francesco Savelli is supported through an Italian State doctoral fellowship.

We are grateful to Alessandra Savelli and Camil Demetrescu for providing valuable introduction to some of the mathematical topics.

References

[1] E. Remolina and B. Kuipers, "Towards a general theory of topological maps," *Artificial Intelligence*, vol. 152, pp. 47–104, 2004.

[2] B. Kuipers, J. Modayil, P. Beeson, M. MacMahon, and F. Savelli, "Local metrical and global topological maps in the hybrid spatial semantic hierarchy," in *Proc. IEEE Int. Conf. on Robotics and Automation*, 2004.

[3] S. Thrun, "Robotic mapping: A survey," in *Exploring Artificial Intelligence in the New Millenium*, G. Lakemeyer and B. Nebel, Eds. Morgan Kaufmann, 2002.

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

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

[6] H. Choset and K. Nagatani, "Topological simultaneous localization and mapping (SLAM): toward exact localization without explicit localization," *IEEE Trans. on Robotics and Automation*, vol. 17, no. 2, 2001.

[7] S. Thrun, "Learning metric-topological maps for indoor mobile robot navigation," *Artificial Intelligence*, vol. 99, no. 1, pp. 21–71, 1998.

[8] E. Fabrizi and A. Saffiotti, "Extracting topology-based maps from gridmaps," in *Proc. IEEE Int. Conf. on Robotics and Automation*, San Francisco, CA, 2000, pp. 2972–2978.

[9] ——, "Augmenting topology-based maps with geometric information," *Robotics and Autonomous Systems*, vol. 40, no. 2, pp. 91–97, 2002.

[10] G. F. Italiano, J. A. L. Poutre, and M. Rauch, "Fully dynamic planarity testing in planar embedded graphs (extended abstract)," in *European Symposium on Algorithms*, 1993, pp. 212–223.

[11] B. J. Kuipers, "The spatial semantic hierarchy," *Artificial Intelligence*, vol. 119, pp. 191–233, 2000.

[12] J. Modayil, P. Beeson, and B. Kuipers, "Using the topological skeleton for scalable global metrical map-building," in *Proc. IEEE/RSJ Int. Conf. on Intelligent Robots and Systems*, Sendai, Japan, 2004.

[13] P. Beeson, M. MacMahon, J. Modayil, J. Provost, F. Savelli, and B. Kuipers, "Exploiting local perceptual models for topological map-building," in *IJCAI-2003 Workshop on Reasoning with Uncertainty in Robotics (RUR-03)*, Acapulco, Mexico, August 2003, pp. 15–22.

[14] J. L. Gross and T. W. Tucker, *Topological Graph Theory*. New York: John Wiley and Sons, 1987.

[15] B. J. Kuipers and T. Levitt, "Navigation and mapping in large scale space," *AI Magazine*, vol. 9, no. 2, pp. 25–43, 1988.

[16] I. M. Rekleitis, V. Dujmović, and G. Dudek, "Efficient topological exploration," in *Proc. IEEE Int. Conf. on Robotics and Automation*, Detroit, USA, May 1999, pp. 676–681.

[17] X. Deng, E. E. Milios, and A. Mirzaian, "Robot map verification of a graph world," *Journal of Combinatorial Optimization*, vol. 5, no. 4, pp. 383–395, 2001.

[18] G. Dudek, P. Freedman, and S. Hadjres, "Using local information in a non-local way for mapping graph-like worlds," in *Proc. of the 13th Int. Joint Conf. on Artificial Intelligence (IJCAI-93)*, 1993, pp. 1639–1645.

[19] R. L. Rivest and R. E. Schapire, "Inference of finite automata using homing sequences," in *Proc. 21st ACM Symposium on Theory of Computing*. ACM, 1989, pp. 411–420.

[20] T. Dean, D. Angluin, K. Basye, S. Engelson, L. Kaelbling, E. Kokkevis, and O. Maron, "Inferring finite automata with stochastic output functions and an application to map learning," *Machine Learning*, vol. 18, no. 1, pp. 81–108, 1995.

[21] M. Bosse, P. Newman, J. Leonard, M. Soika, W. Feiten, and S. Teller, "An Atlas framework for scalable mapping," in *Proc. IEEE Int. Conf. on Robotics and Automation*, 2003.

[22] 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.

[23] B. Kuipers and P. Beeson, "Bootstrap learning for place recognition," in *Proc. of the 18th National Conf. on Artificial Intelligence (AAAI-2002)*, Edmonton, Canada, 2002.