

QUALITATIVE SPATIAL REASONING BASED ON ALGEBRAIC TOPOLOGY

BOI FALTINGS

Swiss Federal Institute of Technology (EPFL)
Switzerland

Abstract. Several formalisms have been proposed for qualitative reasoning about regions and their topological relations in space. These formalisms, based on pairwise relations, do not allow sufficiently powerful inferences to be used for spatial reasoning tasks such as planning a collision-free path. In this paper, I show how considering relations between region *triples*, much more powerful reasoning techniques become possible. I show in particular that in two dimensions, purely topological reasoning is sufficient to compute a *minimal place graph* which represents all minimal and maximal region combinations, as well as all minimal paths between them. I illustrate how this could be applied to motion planning, showing that in spite of its qualitative nature, the formalism is powerful enough to solve problems of practical interest.

1. Introduction

Spatial reasoning problems can often be formulated through topological relations of regions, where the regions are chosen in some way that is suitable for solving the problem at hand. For example, in spatial layout it may be necessary to ensure certain relations between elements. These relations can be expressed as region intersections: for example, A left of B means that A intersects the region “left of B” formed by the points to the left of B. More complex design problems might require the existence of a path with certain characteristics, expressible as constraints on connectivity between regions and their intersections.

Cui, Cohn and Randell (1992) and independently Egenhofer (1991) have proposed techniques for reasoning about pairwise relations of regions in space. Similar in spirit to Allen's logic for temporal intervals (Allen, 1983), these techniques allow completing a partial set of relations with others which follow from them. Often, the inferences are highly

ambiguous, making it difficult to use these logics for practical spatial reasoning problems.

Pairwise relations have been shown to work well in temporal reasoning. This is not surprising, as in the case of one-dimensional regions, Helly's theorem (Chvátal, 1983) allows powerful inferences based only on pairwise relations:

THEOREM 1 *A set R of n convex regions in d -dimensional space has a common intersection if and only if all subsets of $d+1$ regions in R have an intersection.*

Work in temporal reasoning, for example (Vilain, Kautz and Van Beek, 1989), has shown how this theorem implies that a path-consistent temporal constraint network is also globally consistent. Since this result does *not* apply when regions have two or more dimensions, it is not surprising that the inferential power of pairwise relations is much weaker for space than it is for time! Note, however, (Grigni, Papadias and Papadimitriou, 1995) which uses characteristics of planar graphs to extend the inferential power of binary relations.

For spatial reasoning in *two* dimensions, applying Helly's theorem poses two difficulties:

- pairwise relations are not enough, it requires relations between *triples* of regions.
- regions have to be convex. While this is always the case with one-dimensional intervals, it does not hold for many regions in space.

The first problem can be readily solved by working with relations between region *triples*. Fortunately, it turns out that the second problem can be solved as well, since the conditions of Helly's theorem can be considerably weakened (Hadwiger, Debrunner and Klee, 1964):

THEOREM 2 Given a set of regions $R = \{r_1, r_2, \dots, r_n\}$ in d dimensions such that

1. any subset of d regions $\{r_{i_1}, \dots, r_{i_d}\}$ has a single and simply connected intersection, and
2. any subset of $d+1$ regions $\{r_{i_1}, r_{i_2}, \dots, r_{i_{d+1}}\}$ has some intersection, there exists a single, simply connected and nonempty intersection of all regions in R .



Figure 1. Regions in space.

This weakened condition is satisfied in many practical problems. When it is not, it is straightforward to verify this and to decompose regions so that it does become satisfied.

Thus, it is interesting to consider what inferences about 2-dimensional space can be drawn based only on topology and relations between region triples. It turns out that this information is sufficient for surprisingly powerful inferences; it even allows planning collision-free motions in a space of obstacles!

2. Computing Place Graphs Through Region Topologies

I assume a two-dimensional Euclidian space, and consider reasoning about *regions* in this space. I will use lower-case letters to denote regions, for example a , x , r_7 . Upper-case letters will denote sets of regions, e.g. $X = \{a,b,c\}$, and the set of all regions considered is a *universe* U . Throughout this paper, I make the following assumptions:

- regions are simply connected.
- any pair of regions intersects in at most a single connected region.

As input information to my methods, I assume a *region graph*, which is a hypergraph whose nodes are regions, edges connect pairs of regions with a nonempty intersection, and hyper-edges triples of regions with a nonempty intersection. This region graph could be found by systematic geometric intersections, by probing sample points, or it could simply be the way regions are stored in a database. Reasoning will concern *intersections* and *places*, defined as follows:

DEFINITION 1 The *intersection* of a set of regions X , denoted $i(X)$, is defined as the set of points falling simultaneously within all regions in X :

$$i(X) = \{ p \mid (\forall r \in X) \quad p \in r \} = \bigcap_{r \in X} r$$

The place $p(X)$ is the set of points falling exactly within all regions in X :

$$p(X) = \{ p \mid [(\forall r \in X) \quad p \in r] \wedge [(\exists r' \in \mathcal{U} - X) \quad p \in r'] \}$$

Note that $p(X) \subseteq i(X)$. Figure 1 illustrates for $X = \{x, y\}$.

2.1 TWO TOPOLOGICAL INFERENCE MODES

Deciding whether $i(X)$ is empty A first important topological property is whether a set X of regions has a common intersection. In two dimensions, Theorem 2 allows deciding this in time $O(|X|^3)$ by checking that X is a hyper-clique, i.e. that all triples in X are connected by a hyper-arc in the region graph.

If the intersection $i(X)$ is nonempty, so is the intersection of every subset of X . Thus, the set of *all* nonempty intersections can be characterized by the set of *maximal* nonempty intersections. Computing this set corresponds to finding all maximal hyper-cliques in the region graph, an operation which in principle is intractable. However, it turns out that the region graph is a two-dimensional *intersection graph* with properties analogous to *interval graphs* in one dimension. These properties allow for far more efficient graph-theoretic algorithms, described in (Faltings, 1996), which compute the set of all maximal hyper-cliques in time $O(|\mathcal{U}|^2)$, where \mathcal{U} is the universe of regions. They are a generalized version of the famous algorithm of Rose, Tarjan and Leuker(1976) for finding cliques in interval graphs.

Deciding whether $p(X)$ is nonempty

I first define:

DEFINITION 2 *The overlap $o(X)$ of X is the set of points which fall within $i(X)$ and some other regions in the universe:*

$$o(X) = i(X) \cup_{r \text{ u-x}} r$$

Figure 1 shows an example, here $o(\{x,y\}) = i(\{x,y\}) \cup (a \cup b \cup c)$. We then have the theorem:

THEOREM 3 *A place $p(X)$ is nonempty if and only if (1) the intersection $i(X)$ is nonempty, and (2) the overlap $o(X)$ of X is different from $i(X)$ itself.*

Proof:

By definition, the place $p(X)$ contains all points in $i(X)$ but not in $o(X)$. Thus, $p(X)$ can only be empty when either $i(X)$ is empty or $o(X) = i(X)$.

QED

A key idea of this paper is that this condition can often be verified by considering only the *topologies* of $i(X)$ and $o(X)$: if they differ in their topologies, then $i(X)$ must be different from $o(X)$. Mathematically, the topology of a set of points can be characterized by its *homology groups*. In two dimensions, there can be two nonzero homology groups:

- zero-th order H_0 ; its rank is equal to the number of disjoint components, and
- first order H_1 ; its rank equals the number of multiple connections or the number of two-dimensional “holes”.

As an example, $o(\{x,y\})$ in Figure 1 has $|H_0| = 1$ and $|H_1| = 1$. This notation is used to formulate the **topological difference rule**:

THEOREM 4 *If the intersection $i(X)$ is nonempty, and $|H_0(o(X))| \neq |H_0(i(X))|$ or $|H_1(o(X))| > 0$, then the place $p(X)$ is nonempty.*

Proof:

If $o(X)$ is empty, then it is clearly different from $i(X)$. Thus $p(X)$ is nonempty and the rule valid by Theorem 3. Given the underlying assumptions about the regions, the intersection of any set of regions X , $i(X)$ has the topology of a simply connected region, $|H_0(i(X))| = 1$ and $|H_1(i(X))| = 0$. Since identical point sets must have identical homology groups,

any $o(X)$ satisfying the condition cannot be identical to $i(X)$. Thus, again by Theorem 3, $p(X)$ is nonempty.

QED

The rank of the zeroth homology group $|H_0(o(X))|$ can be determined by counting the number of connected components. The rank of the first homology group is then found using the Euler characteristic: $\chi(o(X)) = |H_1(o(X))| - |H_0(o(X))|$, which can be found by a divide-and-conquer algorithm. Details of these algorithms are given in Faltings(1996). They allow computing both homology groups of $o(X)$ for any X .

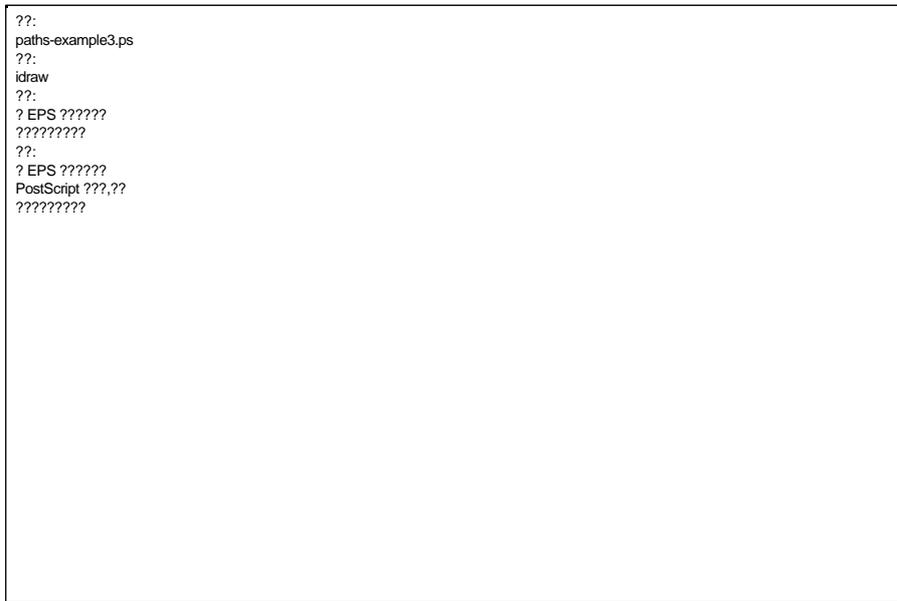


Figure 2. Examples of paths in space.

While this rule gives a *sufficient* condition for $p(X)$, the condition is of course not necessary, and not all non-empty places can be found by this criterion. Next, I define a condition of *minimality* which characterizes the places found by the topological difference rule and turns out to correspond well to the requirements of many practical applications.

2.2 MINIMAL PLACE GRAPHS

I now characterize the places which the topological difference rule can identify:

- *maximal* places are those where $o(X) = \emptyset$, i.e. $|H_0(o(X))| = 0$, $|H_1(o(X))| = 0$.
- *minimal* places are those where $o(X)$ is multiply connected, i.e. $|H_1(o(X))| > 0$.
- *join* places are those where $o(X)$ consists of multiple parts, i.e. $|H_0(o(X))| > 1$.

I then define an adjacency criterion which groups minimal and join places together in a *minimal place graph*. For reasons of space, I do not give proofs of the theorems; these appear in Faltings(1996).

2.3 MINIMAL PLACES

The following definition characterizes *minimal places*:

DEFINITION 3 A place $p(Y)$ is said to be a *direct neighbour* of the place $p(X)$ if and only if for any distance D , there exists a pair of points x and y such that $x \in p(X)$, $y \in p(Y)$ and the distance between x and y is less than D . A place $p(X)$ is *minimal* if and only if it is nonempty and for all places $p(Y)$ which are direct neighbors of $p(X)$, $Y \cap X = \emptyset$.

Minimal places then satisfy the following theorem:

THEOREM 5 Suppose $p(X)$ is a nonempty minimal place in two dimensions among a universe of regions U . Then $|H_1(o(X))| \geq 1$, and $p(X)$ consists of $|H_1(o(X))|$ disjoint nonempty components.

Join places The definition of join places uses the notion of a *monotone* path, defined as follows:

DEFINITION 4 When a point continuously traverses a path, it enters or leaves certain regions. When a path segment only enters regions, it is called *monotone increasing*, when it only leaves regions, *monotone decreasing*. A path is *monotone* if it is either monotone increasing or decreasing.

Join places themselves are then defined as follows:

DEFINITION 5 A place $p(X)$ is a *join* place between two minimal places $p(A)$ and $p(B)$ if and only if:

- there is a monotone increasing path from $p(A)$ to $p(X)$,
- a monotone decreasing path from $p(X)$ to $p(B)$, and

- X is minimal, i.e. there is no other $Y \subset X$ which satisfies these conditions.

In Figure 2, the path from $p(\{d\})$ to $p(\{d,g,h\})$ is monotone increasing, from $p(\{d,g,h\})$ to $p(\{h\})$ monotone decreasing. Since no subset of $p(\{d,g,h\})$ satisfies this condition, it is a join place by Definition 5. Join places satisfy the following theorem:

THEOREM 6 Let $p(X)$ be a join place between two minimal places $p(A)$ and $p(B)$. Then either:

- its overlap $o(X)$ is empty, or
- the rank of the zeroth homology group of $o(X)$, $|H_0(o(X))| > 1$

Theorem 6 shows that for any join place $p(X)$, $|H_0(o(X))| \geq 1$: either $o(X) = \emptyset$, and $p(X)$ is a maximal place found using Theorem 2, or $o(X)$ must be multiply connected. Figure 2 shows examples of join places: those where $|H_0(o(X))| > 0$ in grey, and maximal places (where $o(X) = \emptyset$) in black. Maximal places become join place only when there are in fact no subsets which qualify as join places between the same minimal places.

Minimal + join places = minimal place graph By Theorems 5 and 6, the set of all minimal and join places is found as the set of maximal places plus all nonempty intersections which satisfy the topological difference rule. I now define an adjacency criterion which joins them together in a *minimal place graph*.

DEFINITION 6 A place $p(X)$ is *directly adjacent* to $p(Y)$ if and only if:

1. $X \cap Y = Z$ ($Z(X \cap Z = Y)$) ($X \cap Y = Z$ ($Z(X \cap Z = Y)$))

where all Z must be such that $p(Z)$ is a minimal or join place, and 2. one of $p(X)$, $p(Y)$ is a join place, and the other is a minimal place.

We then have the following:

LEMMA 1 Let $p(A)$ and $p(B)$ be two directly adjacent places in the minimal place graph. Then there is a monotone path from a point in $p(A)$ to a point in $p(B)$.

When the universe U of regions does not cover the entire space, the empty set $p(\{\})$ is also a nonempty place. The above rule cannot be used to determine its adjacencies. Instead, the empty place is adjacent to all join places which have only a single neighbour.

Here are some adjacencies for the example of Figure 2:

$p(\{\})$ $p(\{a,c\})$ $p(\{c\})$
 $p(\{c\})$ $p(\{c,d\})$ $p(\{d\})$
 $p(\{d\})$ $p(\{d,g,h\})$ $p(\{h\})$

2.4 THE MINIMAL PLACE GRAPH AS A SOUND AND COMPLETE REPRESENTATION OF REALITY

While the minimal place graph does not contain all physically realizable places, it is sufficient for many purposes. We now show that the minimal place graph as a representation of the structure of the space is:

- complete: every path and place in real space has a corresponding path or place, respectively, in the minimal place graph, and
- sound: every path and place in the minimal place graph is physically realizable.

First, every physically realizable place has some equivalent in the minimal place graph:

THEOREM 7 *Given any nonempty place $p(X)$, there exists a place $p(Y)$ in the minimal place graph such that $Y \subseteq X$ and furthermore there is a monotone decreasing path from a point in $p(X)$ to a point in $p(Y)$.*

The minimal place graph contains only physically realizable places by construction. The following two theorems show that the minimal place graph correctly mirrors the connectivity of minimal places:

THEOREM 8 Let $X = (p(X_0), p(X_1), \dots, p(X_k))$ be the sequence of all places traversed by a physically realizable path for connecting a point in $p(X_0)$, with a point in $p(X_k)$. (This means that all successive places $p(X_i)$ must be direct neighbours, i.e. X_i and X_{i+1} differ by a single region. Then there exists an equivalent sequence $Y = (p(Y_0), p(Y_1), \dots, p(Y_l))$ such that:

- $X_0 \supseteq Y_0, X_k \supseteq Y_l$
- all Y_i are subset of some X_j
- all successive places $p(Y_i)$ are minimally adjacent
- all $p(Y_i)$ are either minimal or join places

THEOREM 9 Let $X = (p(X_0), p(X_1), \dots, p(X_k))$ be a path in the minimal place graph, where all $p(X_i)$ are simply connected regions. Then there exists a path through points (x_1, x_2, \dots, x_k) in physical space such that:

1. $x_i \in p(X_i)$

2. there exist monotone paths between x_i and x_{i+1} passing only through regions in X_0, X_1, \dots, X_{ki} , but not crossing any other regions in the universe.

Theorem 7 shows that there is a monotone decreasing path from an arbitrary point to a point in a minimal place. Consequently, for any path P in physical space, there is a path P' starting and ending in minimal places which does not traverse any regions not traversed by P as well. Theorem 8 then shows that any such path P' has an equivalent in the minimal place graph. Together, the two theorems imply completeness: every physically realizable path P has an equivalent P' in the minimal place graph. Furthermore, its equivalent is qualitatively at least as good as any of the equivalent physical paths it represents in the sense that it does not traverse more regions than the path in the full graph¹.

Theorem 9 shows that whenever there is a path in the minimal place graph passing only through simply connected places, then there also exists a corresponding path in physical space. It establishes soundness: all paths in the minimal place graph are also physically realizable. However, soundness only holds under the restriction that the places must be simply connected; this is the question we address next.

Figure 2 shows several examples of paths. The two paths shown with solid arrows show two ways of connecting a point in $p(\{c\})$ with a point in $p(\{h\})$, either by passing through the interior or by passing through the empty space outside the regions. Note that the path shown by the dashed arrow is not explicitly represented in the place graph. This is because it does not require passing through any other regions than those of the origin and destination. If such paths should be ruled out, then the empty space must be covered with explicit regions (Figure 3 uses this). In the example of Figure 2, this would make all single regions become explicit minimal places in the place graph.

Region Topology Theorem 9 guarantees a mapping from paths in the minimal place graph to paths in physical space only under the condition that all places are *simply connected*.

¹

Although by quantitative criteria, it may be longer.

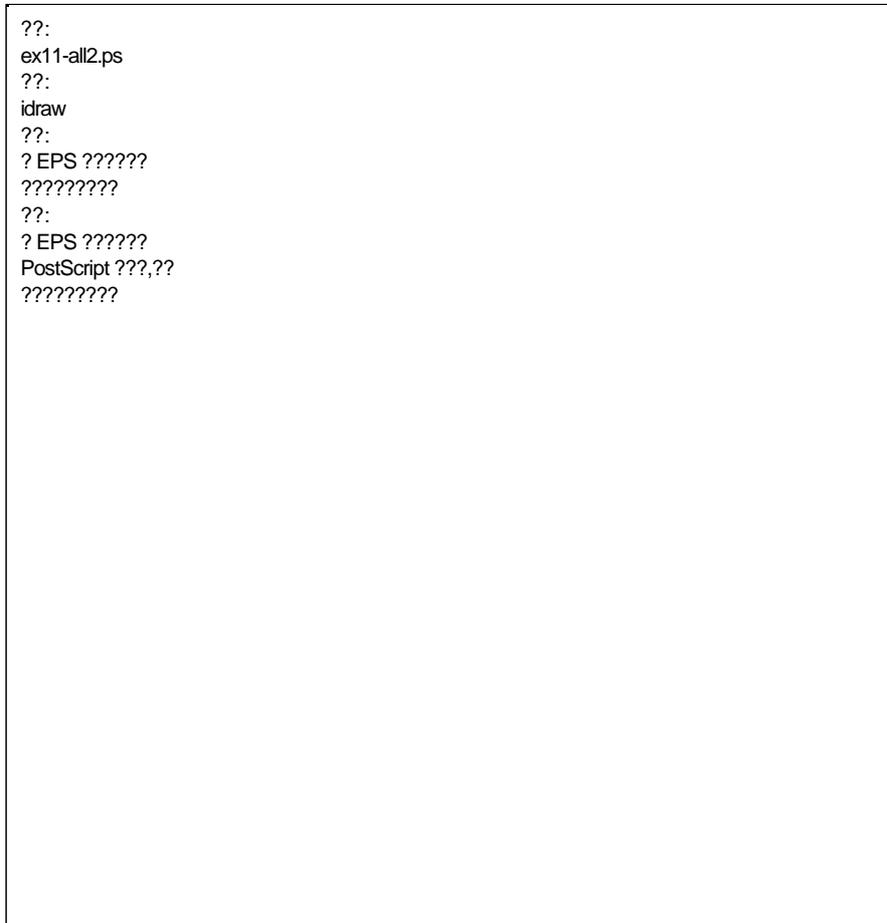


Figure 3. Input representation of an instance of the piano movers problem.

When place consist of several parts, paths passing through the place do not necessarily reflect physically possible paths. Fortunately, this problem is easy to detect and correct. It is detected when the topology of the overlap of X is more than doubly connected, i.e. $|H_1(o(X))| \geq 2$. It can be corrected by adding or subdividing regions such that the different components of the place fall within different regions and thus become different places.

3. Example: The Piano Movers Problem

As an example of an application of the formalism I developed, consider the *piano movers problem* (Schwartz & Yap, 1987), a problem which has

drawn much attention in robotics and spatial reasoning. In the piano movers problem, the goal is to find a path for moving a single rigid *moving object* from an initial to a final position such that it does not collide with any of the fixed and rigid obstacles.

```

??:
ex11-all2.ps
??:
idraw
??:
? EPS ??????
?????????
??:
? EPS ??????
PostScript ???,??
?????????

```

Figure 3. Small input representation of an instance of the piano movers problem.

In a qualitative version of the problem, positions of the moving object are given by regions within which all points are considered equivalent. As an example, consider the situation shown in Figure 3. Here, convex regions are used to model a moving object, consisting of regions x and y , a set of obstacles (hatched), and a covering of the empty space between them. Regions making up the obstacles and free space are numbered sequentially. The scenario is enclosed by a frame which serves to rule out paths through empty space, as mentioned earlier.

A configuration is a particular position and orientation of the moving object and defines a point in a configuration space (Lozano-Perez and Wesley, 1979). I define a qualitative configuration of the moving object by the combination of regions which the moving object overlaps. I use the notation a/b , where a denotes the part of the moving object and b the region it overlaps, e.g. configuration $A=\{x/8,y/8,y/9\}$ and $B=\{x/10,y/10,y/11\}$ and $B=\{x/1,y/1,y/5,x/10,y/10\}$. Thus, positions of the moving objects are represented in purely topological terms, without recourse to any coordinate system.

Each possible overlap between a part of the moving object and one of the other regions defines a region of points in configuration space which I call a *c-region*. Because of the convexity of object parts, it is possible to show:

THEOREM 10 *Every c-region formed by two convex pieces or cavities A and B is a convex region.*

Thus, it is possible to apply the topological inference techniques described in this paper to compute a minimal place graph representing qualitatively the different paths and allowing us to decide whether or not there is in fact a path between a given pair of positions. Figure 4 shows the legal part of the minimal place graph for this example. It qualitatively represents all legal motions of the moving object. A qualitative solution to the piano movers problem can be given by first mapping the initial and final positions to the minimal places to which they belong, and then finding the qualitative path between initial and final place by searching in the graph. For example, a path between configurations A and B (Figure 3) can be found as follows. First, I map to the minimal places: $A=\{x/8,y/8,y/9\}$ is already minimal, $B=\{x/10,y/10,y/11\}$ is mapped to $\{x/10,y/10\}$. Next, I find a path in the graph between the two minimal places, in this case the path $\{x/8,y/8,y/9\} ? \{X/8,y/8,X/9,y/9,y/10\} ? \{x/9,y/9,y/10\} ? \{x/9,x/10,y/9,y/10\} ? \{x/10,y/10\}$.

I have implemented a prototype which demonstrates the topological reasoning techniques on the piano movers problem for two-dimensional objects. The input to the program is given in the form of three collections of convex bitmaps, representing the parts of the fixed objects, the moving object, and the cavities. A preprocessor uses these bitmaps to determine all possible simultaneous overlaps of 3 pairs of parts. Topological reasoning then determines the minimal place graph and thus possible paths. The implementation shows that it is possible to solve even such complex

planning problems without *any* analytical representations of shapes, without any fixed resolution limits or other approximations.

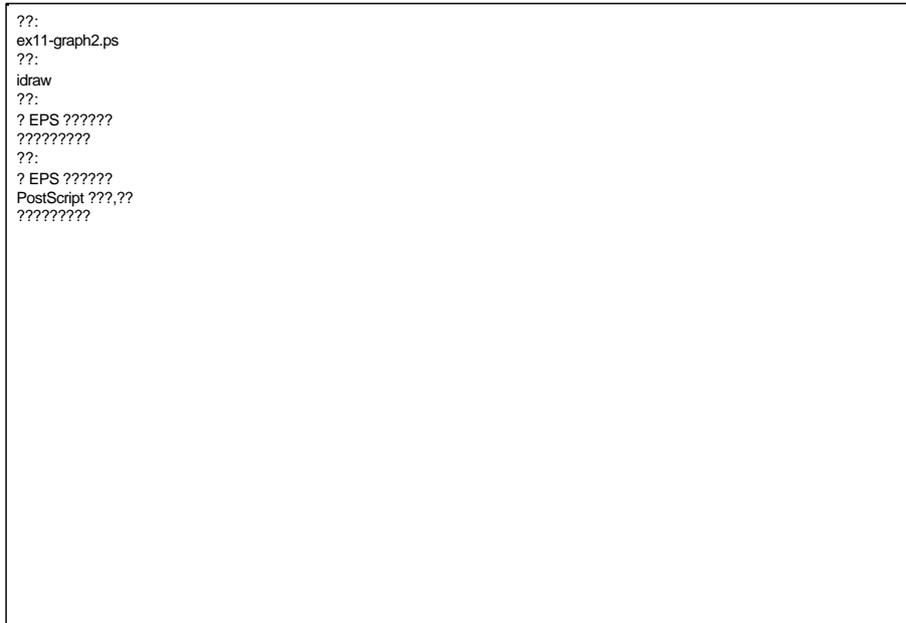


Figure 4. The legal part of the feasible skeleton for the example. All places containing overlaps (obstacles) have been omitted. The numbering of the regions refers to Figure 4. Minimal places are shown in black, join places in white.

4. Conclusions

Qualitative spatial reasoning is an important and largely unsolved problem. Concepts of topology are promising for two reasons. First, they do not require any coordinate system. In other techniques, such coordinate systems must often be imposed arbitrarily and influence the result of reasoning in an equally arbitrary way. Second, they allow treating curved objects without greatly increasing the complexity. Natural shapes are usually curved, so this is an important advantage.

With the exception of (Grigni, Papadias and Papadimitriou, 1995), existing work on topological reasoning using binary relations does not exploit the fact that regions are *spatial*; inferences are valid for any sets. Thus, it is not surprising that their inferential power is rather weak. In this paper, I have shown that by using *ternary* instead of binary relations, much

more powerful inference becomes possible. Even the piano movers problem, one of the hardest spatial reasoning problems, can be solved! There are many other applications of topological inference, especially in the area of geographical information systems. Another promising area is the combination of qualitative spatial reasoning to constrain the result of quantitative geometric computation.

References

- Allen, J.F.: 1983, Maintaining knowledge about temporal intervals, *Communications of the ACM*, **26**(11): 832-843.
- Cui, Z., Cohn, A.G. and Randell, D.A.: 1992, Qualitative simulation based on a logical formalism of space and time, *Proceedings of the 10th National Conference of the AAAI*, San Jose.
- Chvátal, V.: 1983, *Linear Programming*, W. H. Freeman.
- Egenhofer, M.J.: 1991, Reasoning about binary topological relations, in O. Gunther and H.J. Schek (eds), *Advances in Spatial Databases, SSD'91 Proceedings*, Springer-Verlag, pp. 143-160.
- Faltings, B.: 1996, Spatial reasoning in two dimensions using topology, *Technical Report EPFL-DI 96/217*, Swiss Federal Institute of Technology, Lausanne.
- Grigni, M., Papadias, D. and Papadimitriou, C.: 1995, Topological inference, *Proceedings of the 14th International Joint Conference on Artificial Intelligence*, Morgan-Kaufman.
- Hadwiger, H., Debrunner, H. and Klee, V.: 1964, *Combinatorial Geometry in the Plane*, Holt, Rinehart and Winston, New York.
- Lozano-Perez, T. and Wesley, M.: 1979, An algorithm for planning collision-free paths among polyhedral obstacles, *Comm. of the ACM*, **22**.
- Schwartz, J.T. and Yap, C.K.: 1987, *Advances in Robotics, Vol. 1: Algorithmic and Geometric Aspects of Robotics*, Erlbaum, Hillsdale, N.J.
- Rose, D.J., Tarjan, R.E. and Leuker, G.S.: 1976, Algorithmic aspects of Vertex elimination on graphs, *SIAM J. Comput.* **5**: 266-283.
- Vilain, M., Kautz, H. and van Beek, P.: 1989, Constraint propagation algorithms for temporal reasoning: A revised report, in D. S. Weld and J. de Kleer (eds), *Readings in Qualitative Reasoning about Physical Systems*, Morgan Kaufmann, pp. 373-381.