Topological Inference

Michelangelo Grigni* and Dimitris Papadias† and Christos Papadimitriou‡
Computer Science and Engineering
University of California, San Diego
9500 Gilman Drive
La Jolla, CA 92093–0114

Abstract

Geographical database systems deal with certain basic topological relations such as "A overlaps B" and "B contains C" between simply connected regions in the plane. It is of great interest to make sound inferences from elementary statements of this form. This problem has been identified extensively in the recent literature, but very limited progress has been made towards addressing the considerable technical difficulties involved. In this paper we study the computational problems involved in developing such an inference system. We point out that the problem has two distinct components that interact in rather complex ways: *relational consistency*, and *planarity*. We develop polynomial-time algorithms for several important special cases, and prove almost all the others to be NP-hard.

1 Introduction

Suppose that you are told that a simply connected planar region A overlaps another region B, and that one of the regions A and C contains the other but you don't know which. What can you infer about B and C? For example, could they be disjoint?

Such questions are of great interest for developing intelligent inference engines for geographic database systems. It has been recently pointed out [Egenhofer,1991] that there are eight fundamental relations that can hold between two planar regions: "overlaps," "disjoint," "inside," "contains," "meets" (overlaps only at the boundary), "covers" (contains but also shares some boundary), "covered by" (the inverse of "covers") and "equal" (see Figure 1). We call these relations *high resolution case*; they are the only relations that can be defined by considering intersections of two regions, their boundaries, and their complements [Egenhofer,1991].

Now, any three planar regions cannot stand in arbitrary relation with respect to each other; for example, if A is inside B and B meets C, then A and C must be disjoint. The complete table of such one-step inferences was derived in [Egenhofer,1991; Smith and Park,1992] (see Table 1; notice that this is an extension —as we shall see, a surprisingly subtle one— of Allen's classical work on temporal intervals [Allen,1983]).

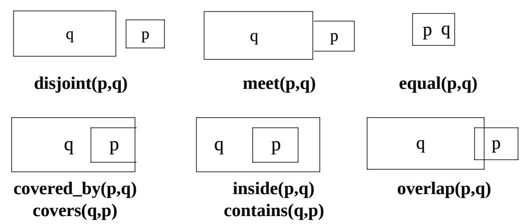

Figure 1: topological relations (high resolution case)

In some cases the refinement provided by the high resolution relations is not needed. In a cadastral application, for instance, the difference between "inside" and "covered by" may not be important. Consider the query "find all land parcels in a given area." The land parcels of the result should be inside or should be covered by the area. In this paper we also focus on the (possibly more useful) case in which there is no differentiation between "meets" and "overlaps" (they are both called "overlap") or between "covers" and "contains" (they are both called "contains"), or their inverses (they are both called "inside"). It can be shown that these are the only relations that are relevant if one considers intersections of two objects and their complements (but not of their boundaries). The one-step inferences of these *medium resolution* relations are given in Table 2.

In addition, we consider a sub-case of medium resolution in which the relation "overlap" is not permitted (instead objects can only meet). This situation arises often in geographic applications where geographic regions and administrative subdivisions obviously can only meet or contain one another, but cannot overlap. The one-step inferences in this case are shown in Table 3.

We also consider the even coarser situation in which there are only two possibilities of interest: "disjoint" and "overlap" (Table 4). Overlap in this case is just the negation of disjoint (it means the two objects have points in common). These two relations are important

*Work supported by NSF grant No. DMS-9206251.
†Work supported by NSF-IRI 9221276.
‡Work supported by NSF grant No. CCR-9301031.

	disjoint d	meet m	equal e	inside i	coveredby cb	contains ct	covers cv	overlap o
disjoint	d ∨ m ∨ e ∨ i ∨ cb ∨ ct ∨ cv ∨ o	d ∨ m ∨ i ∨ cb ∨ o	d	d ∨ m ∨ i ∨ cb ∨ o	d ∨ m ∨ i ∨ cb ∨ o	d	d	d ∨ m ∨ i ∨ cb ∨ o
meet	d ∨ m ∨ ct ∨ cv ∨ o	d ∨ m ∨ e ∨ cb ∨ cv ∨ o	m	i ∨ cb ∨ o	m ∨ i ∨ cb ∨ o	d	d ∨ m	d ∨ m ∨ i ∨ cb ∨ o
equal	d	m	e	i	cb	ct	cv	o
inside	d	d	i	i	i	d ∨ m ∨ e ∨ i ∨ cb ∨ ct ∨ cv ∨ o	d ∨ m ∨ i ∨ cb ∨ o	d ∨ m ∨ i ∨ cb ∨ o
covered by	d	d ∨ m	cb	i	i ∨ cb	d ∨ m ∨ ct ∨ cv ∨ o	d ∨ m ∨ e ∨ cb ∨ cv ∨ o	d ∨ m ∨ i ∨ cb ∨ o
contains	d ∨ m ∨ ct ∨ cv ∨ o	ct ∨ cv ∨ o	ct	e ∨ i ∨ cb ∨ ct ∨ cv ∨ o	ct ∨ cv ∨ o	ct	ct	ct ∨ cv ∨ o
covers	d ∨ m ∨ ct ∨ cv ∨ o	m ∨ ct ∨ cv ∨ o	cv	i∨ cb ∨ o	e ∨ cb ∨ cv ∨ o	ct	ct ∨ cb	ct ∨ cv ∨ o
overlap	d ∨ m ∨ ct ∨ cv ∨ o	d ∨ m ∨ ct ∨ cv ∨ o	o	i∨ cb ∨ o	i∨ cb ∨ o	d ∨ m ∨ ct ∨ cv ∨ o	d ∨ m ∨ ct ∨ cv ∨ o	d ∨ m ∨ e ∨ i ∨ cb ∨ ct ∨ cv ∨ o

Table 1: high resolution one-step inferences

	disjoint	equal	inside	contains	overlap
disjoint	d ∨ e ∨ i ∨ ct ∨ o	d	d ∨ i ∨ o	d	d ∨ i ∨ o
equal	d	e	i	ct	o
inside	d	i	i	d ∨ e ∨ i ∨ ct ∨ o	d ∨ i ∨ o
contains	d ∨ ct ∨ o	ct	ct ∨ e ∨ i ∨ o	ct	ct ∨ o
overlap	d ∨ ct ∨ o	o	o ∨ i	d ∨ ct ∨ o	d ∨ e ∨ i ∨ ct ∨ o

Table 2: medium resolution one-step inferences

	disjoint	equal	inside	contains	meet
disjoint	d ∨ e ∨ i ∨ ct ∨ m	d	d ∨ i ∨ m	d	d ∨ i ∨ m
equal	d	e	i	ct	m
inside	d	i	i	d ∨ e ∨ i ∨ ct ∨ m	d ∨ m
contains	d ∨ ct ∨ m	ct	ct ∨ e ∨ i	ct	ct ∨ m
meet	d ∨ ct ∨ m	m	m ∨ i	d ∨ m	d ∨ e ∨ i ∨ ct ∨ m

Table 3: one-step inferences without overlap

in spatial access methods where tree structures, such as R-trees, are used to efficiently answer queries of the form "find all objects that overlap with object A" [Papadias *et al.*,1995].

	disjoint	**overlap**
disjoint	d ∨ o	d ∨ o
overlap	d ∨ o	d ∨ o

Table 4: low resolution one-step inferences

Therefore we can distinguish three levels of qualitative resolution (the medium level consists of two sub-cases). The choice of topological relations depends on the resolution requirements of the specific application domain. The three levels are illustrated in Figure 2.

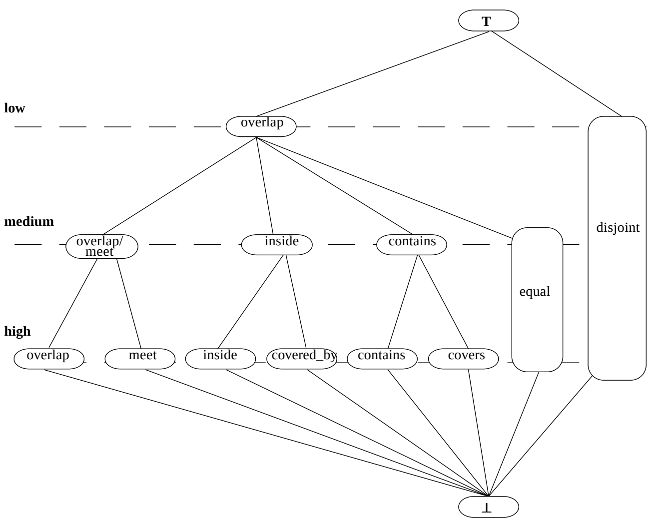

Figure 2: levels of topological resolution

We are interested in deriving inferences involving topological relations such as the above. Instead of considering the inference problem directly, we shall focus on the corresponding *satisfiability problem*, that is, determining whether or not a Boolean combination of statements of the form "A meets B," "B overlaps C" etc. can hold for certain planar regions A, B, C, etc. If the expression is unsatisfiable, we should be able to conclude so; if it is satisfiable, we should be able to come up with *actual planar regions that satisfy it.*

The paper is organized as follows: Section 2 describes in detail the problem and its sub-cases. Section 3 summarizes the computational complextiy results, classifying each problem in a complexity class. Section 4 describes proof sketches and Section 5 concludes with comments about future work.

2 Problem Description

We shall only consider Boolean expressions that are conjunctions of clauses, where each clause is of the form ("A meets B" or "A overlaps B" or "A inside B"), that is, the conjunction of one or more atomic statements, *all involving the same pair of objects*; furthermore, without loss of generality, there is exactly one clause involving each pair of objects (if there are two or more, then this is equivalent to the disjunction of the relations that are common to all clauses; if there is none, then we implicitly have the *full clause*, the disjunction of all possible relations at the present resolution). We call such expressions *topological expressions;* the generalization to arbitrary Boolean expressions in CNF involves no additional difficulty, and would change our results very little (see Entry 5 in Section 4) but also appears to be of no use.

There are two important special cases of topological expressions that are of interest: In the *explicit* case, all clauses are singletons (the relation for each pair of objects is known). In the *conjunctive* case we only have clauses that are either singletons or full (the relation for each pair of objects is either explicit or unknown). This situation arises in geographic applications where the relation between objects in the same map is known, but not explicit information is given about objects in diffferent maps.

We consider two example topological expressions involving objects A, B, C, and D in medium resolution.

$$\text{(A overlaps B)} \wedge \text{(B contains C)} \wedge \text{(A inside D)} \\ \wedge \text{(B disjoint D)} \wedge \text{(C disjoint D)}$$

This first expression is conjunctive but not explicit: there are no ∨'s, but there is a pair (A, C) whose relation is unspecified. The pair is only related by an implicit full clause (A overlaps C ∨ A disjoint C ∨ A equal C ∨ A inside C ∨ A contains C).

$$\text{(A overlaps B} \vee \text{A equal B)} \wedge \text{(B contains C)} \\ \wedge \text{(A inside D} \vee \text{A contains D)}$$

This second expression is not conjunctive (there are clauses that are not full, but have at least two disjuncts). The second is satisfiable (by three cocentric circles, starting from the outermost, A=B, C, and D), whereas the first is not (A, B, and D contradict first-column, last-row entry of Table 2).

There are two distinct kinds of reasons why a topological expression may be unsatisfiable. First, it may contradict the relational consistency as expressed by the inference tables (Tables 1 through 4). This aspect of the satisfiability problem is in effect a *constraint satisfaction problem*, and has been studied as such.

The second aspect is more subtle, and had escaped the researchers in this area[1]. A set of relations may be consistent, and still there may be no planar regions that realize it because of reasons related to *planarity.* For example, it is well-known that the complete graph with five nodes is non-planar (see Figure 3). Assume that we are given the objects X1, X2, X3, X4, X5, Y1, ..., Y9 related as shown (all objects are disjoint except

[1]It is easy to understand why this important point had not been noticed before. Topological reasoning is a two-dimensional extension of the classical work of Allen [Allen,1983] on reasoning about temporal intervals. In the case of intervals, however, constraint satisfaction is enough; in the case of planar regions it is not.

that X1 overlaps Y1, Y2, Y3, and Y9, and so on for the other nodes and edges). In addition we are given that Y10 overlaps X1 and X4 and is disjoint with every other object. Then relational consistency according to Tables 1-4 will reveal no contradiction. For instance, in the simplest case of Table 4 any topological relation between a pair of objects is permitted regardless of the relations between the other pairs of objects. However, there is no way to realize this set of relations by a set of regions in the plane, without having two of the Y objects overlap. If we try to insert Y10 in the partial realization of Figure 3, then Y10 must overlap with at least one other Y object (Y8 illustrated).

Therefore, there are some subtle geometric constraints that must be satisfied, besides the relational consistency ones; notice the analogue with the scene recognition problem [Kirousis and Papadimitriou,1988], where, besides the relational consistency constraints of labels of the edges of a scene ("convex," "concave," "boundary"), there are subtle additional geometric constraints (the planar regions must be realizable in terms of actual slopes and heights).

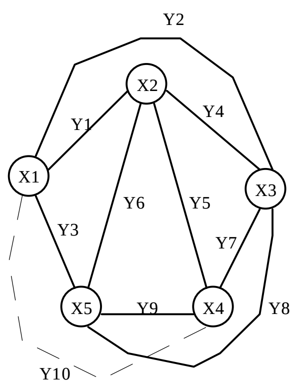

Figure 3: planarity obstructions

We start by addressing the constraint satisfaction aspects of the problem. In a constraint satisfaction problem we are given variables $x_1, \ldots, x_n$ taking values in corresponding finite domains $D_1, \ldots, D_n$. We are also given constraints on subsets of the variables: a constraint $R_{i_1,\ldots,i_k}$ is a k-ary predicate on the values of $x_{i_1}, \ldots, x_{i_k}$ (we will only consider ternary constraints). The problem is to find an assignment of values to the variables such that all the constraints are satisfied. Constraint satisfaction problems are typically NP-complete, although special classes can be solved in polynomial time.

In the constraint satisfaction problems arising in connection to topological inference, the variables are *pairs of distinct objects*; all domains are the subsets of the set of eight topological relations in Table 1 (or the five in Tables 2 and 3, or the two in Table 4), as dictated by the clause corresponding to the pair of objects; and for each triple (i, j, k) of objects we have a constraint, namely, that the value of the pair i, j, the value of the pair j, k, and the value of the pair i, k must be related as in Table 1 (or 2, or 3, or 4).

For example, the topological expression

$$(\text{A overlaps B} \vee \text{A equal B}) \wedge (\text{B contains C}) \\ \wedge (\text{A inside D} \vee \text{A contains D})$$

is expressed by six variables (all unordered pairs of objects from A, B, C, D). The domain of variable AB is the set {overlaps, equal}, the domain of BC is {contains}, and the domain of the unrestricted pair AC is all five medium resolution relations. There are twenty-four ternary constraints (one for each ordered triple), expressed by Table 2. A satisfying assignment assigns equal to AB, contains to AC, inside to AD, contains to BC, inside to BD, and inside to CD.

3 Summary of Results

The satisfiability problem for topological expressions has several special cases and subproblems along three dimensions:

1. **Generality of the Topological Expression.** There are three cases of interest, *unrestricted topological expressions*, *conjunctive topological expressions*, and *explicit topological expressions.* Obviously, their complexity is non-increasing in this order.
2. **Resolution of the Topological Relations.** We consider four cases: In the finest level we have the eight topological relations of Figure 1; this is the *high resolution case.* In the next level we overlook the difference between contains and covers, and between overlaps and meets; this is the "medium resolution case." In the coarsest level, we only have the relations "disjoint" and "overlap" (this is the *low resolution case.* Finally, we also consider the special case of the second level where we can have no "overlap," but "meet" (this is the *no-overlap* case). There is no *a priori* dominance of the complexities of these cases.
3. **Notion of Satisfiability.** We consider two notions of satisfiability: *Relational consistency* means that there is a global choice of a disjunct from the clause of each pair of objects that, so that all triples of choices are consistent with the table of the present resolution. As the example in Figure 3 indicates, this concept is sound but not complete (its answer to the satisfiability question may be a false positive, but never a false negative). The full form of satisfiability is called *realizability:* A topological expression on a set of objects is said to be realizable if it has a *planar model*; that is, if there is a set of simply connected planar regions, one for each object, any two of which are related by a topological relation that is a disjunct of the corresponding clause.

Tables 5 and 6 summarize our results. The rows refer to the level of generality of the topological expressions. The columns correspond to the resolution of the topological relations. The first table indicates the complexity of the relational consistency problem, while the second table the complexity of the realizability problem. "P" means that the problem is solvable in polynomial time; "NP-h" that it is NP-hard.

	high	medium	low	no overlap
unrestricted	NP-h^1	NP-h^3	P^5	NP-h^7
conjunctive	P^9	P^{11}	P^{13}	P^{15}
explicit	P^{17}	P^{19}	P^{21}	P^{23}

Table 5: Complexity of relational consistency.

	high	medium	low	no overlap
unrestricted	NP-h^2	NP-h^4	NP-h^6	NP-h^8
conjunctive	NP-h^{10}	NP-h^{12}	NP-h^{14}	?^{16}
explicit	NP-h^{18}	NP-h^{20}	NP-h^{22}	P^{24}

Table 6: Complexity of realizability.

According to the above tables, the only open question is entry 16, the realizability problem for the no-overlap conjunctive case. The corresponding relational consistency problem (that is, telling whether there is a relationally consistent solution) is in P (entry 15); furthermore, the problem of telling whether an explicit expression (a solution, that is) is realizable is also in P. However, the combined problem is not at all clearly in P (there may be too many solutions to consider); there is a parallel in [Kirousis and Papadimitriou,1988].

4 Proof Sketches

In this section we sketch proofs of the results. We begin with the polynomial time algorithms.

Entry 11. Given a conjunctive medium resolution formula ϕ, we first require that it pass the path-consistency algorithm with respect to Table 2 (running time $O(n^3)$ [Mackworth and Freuder,1985]). Assuming this, augment ϕ with all explicit (non-disjunctive) clauses derived by path-consistency. Finally assert "disjoint" for all remaining pairs.

To see that this works, consider ϕ as augmented. The "equals" relations define an equivalence relation on the regions, the "inside" and "contains" relations define an irreflexive partial order on the equivalence classes, the "disjoint" relations are closed downward with respect to the order, and the "overlaps" relations are closed upward. We now construct a *set model* of the regions. Let U be the universe consisting of all regions X together with all pairs YZ such that "Y overlaps Z" is asserted in ϕ. Let "$\leq$" denote "inside or equals." Model region A by a subset of U: $S(A) = \{X : X \leq A\} \cup \{YZ : Y \leq A \lor Z \leq A\}$. Note Table 2 is consistent with the relations of *any* system of sets. Furthermore, these sets agree with ϕ, hence they tell us how to complete ϕ to explicit form. Now observe that for a pair AB not already related in ϕ, $S(A)$ and $S(B)$ are disjoint.

Entry 9. The consistency proof now uses a *pair* of sets $S_1(A) \subset S_2(A)$ to model each region. We check path consistency, and assert "disjoint" or "meet" for undetermined pairs. Details omitted.

Entry 15. This is nearly identical to Entry 11, replacing "overlaps" with "meets," again asserting "disjoint" for pairs not determined by path-consistency. Whenever we have $S(A) \subseteq S(B) \cap S(C)$ in this model, either $S(B)$ and $S(C)$ are equal or one contains the other. Such sets are consistent with Table 3.

Entries 17, 19, 23. Path-consistency suffices here. However, path consistency is *not* complete (even for relational consistency) in the unrestricted case. That is, there are examples in which the path consistency algorithm will halt without identifying an inconsistency, and still the expression is unsatisfiable (simple examples can be found using our NP-completeness construction below).

Entry 24. This is basically a planarity problem, in a graph that combines the "contains" and the "meets" relations. The only problem is that four or more regions may meet pairwise at a point, resulting in large cliques (graphs with large cliques are, of course, nonplanar). This is circumvented by a pre-processing phase that identifies all maximal cliques of size four or more, and replaces them by a new vertex connected to each of the vertices in the clique (intuitively, the new vertex stands for the common point). The maximal cliques of a graph can be identified in time polynomial in the nodes *and the number of maximal cliques*, which must be polynomial for any realization.

Entries 5, 13, 21. Since Table 4 imposes no real constraint in low resolution, all topological expressions are satisfiable. Incidentally, if we allow general Boolean combinations of relational statements instead of topological relations (that is, clauses that involve statements about more than one pair of objects), then this entry becomes NP-hard: We can simulate Boolean satisfiability by having an object for each variable, and also another object 0, and replacing all instances of variable x with "0 disjoint x" and the negation of x with "0 overlap x". Entry 5 is the only result in our tables that would be changed by allowing general Boolean expressions.

Next we sketch the NP-hardness results.

Entries 1, 3, and 7. The proof for Entry 3 is by a reduction from the "not-all-equal satisfiability" problem. We notice that if A and B contain one another but we do not know which, the same for B and C, and if we know that A and C overlap, then we can conclude that either both A contains B and C contains B, or the opposite. Hence, if "A contains B" stands for "x is true" and "A inside B" for "x is false," this observation allows us to "propagate the truth values" of the variables. The "not-all-equal" clauses are simulated by a triangle of "contains-or-inside" relations, that cannot all three go clockwise, or all three counterclockwise, because this would violate transitivity. The same reduction works for Entry 1, while for entry 7 a more complex "value propagation gadget" is needed.

Entries 10, 12, 14, 18, 20, and 22. The easiest problem among these is Entry 22, and it turns out to be NP-hard because of a result of [Kratochvíl,1991]: It is NP-hard to tell whether a given graph is a string graph; that is, whether there exists a set of curves, one for each node, such that for any two nodes the corresponding curves intersect if and only if the nodes are adjacent. Furthermore the result of [Kratochvíl and Matoušek,1991] that some string graphs require exponential size realizations carries over to the case of regions as well, so it is not clear whether any of these problems is in NP.

5 Discussion and Further Questions

The qualitative representation and processing of spatial knowledge has recently gained much attention in spatial databases (see for example [Papadias and Sellis,1994]), because very often in geographic applications we need to handle spatial relations such as, "disjoint", "overlap" and "north". In this paper we assume that we are given a database of objects with their interrelationships explicitly represented. The data may be incomplete (we may have no information for some pairs of objects), indefinite (more than one relation between the same pair of objects is possible) or inaccurate (the relations may lead to non-realizable configurations of objects). Such problems arise often in spatial databases and geographic information systems (GIS) where data from various sources and of variable quality are incorporated in the same system. This paper discusses algorithms that can be used to infer the relation between the pairs of objects for which the spatial relation is not known and to prune the impossible relations.

The proposed methods refer to configurations of arbitrary objects. On the other hand geographic databases contain particular objects whose shape and size further restrict the allowable topological relations. An interesting extension of the work in this paper would be to study the topological inference problem when we know something about the regions (they are all convex, or circles, or rectangles, or even given polygons). Despite the fact that satisfiability for particular objects would require further processing that takes into account the sizes and the shapes of the objects, our method can be used as a pre-processing step in commercial GIS's that currently do not involve any inference or satisfiability mechanisms.

In this paper we have demonstrated that several practical problems can be solved by reasonably sophisticated polynomial algorithms, while for the rest we have shown that only non-polynomial solutions may exist for the general case. Even for the NP-hard cases, we suspect that *heuristics* will be of great use. Naturally, standard or specialized constraint satisfaction heuristics can solve the relational consistency subproblem. As for the planar realizability subproblem, which we showed almost everywhere NP-hard, there is hope. Our counterexample showing that planarity is an additional constraint (Figure 3) involves *fifteen* regions; in the low-resolution case, the smallest non-realizable example involves *twelve* regions. One could devise heuristics that have an excellent chance of finding a planar model, if one exists.

Finally, further work can be done for other types of spatial relations, such as cardinal directions (e.g., "north," "northeast") and distance relations ("near", "far") and applications that involve several kinds of qualitative and quantitative spatial constraints ("10 km to the north", "disjoint but near"). When choosing a set of relations to model, a reasonable criterion is to choose relations where at least the conjunctive relational consistency problem is solvable in polynomial time.

References

[Allen, 1983] J. F. Allen Maintaining Knowledge about temporal intervals *Communications of the ACM*, 26(11):832–843, 1983.

[Egenhofer, 1991] M. J. Egenhofer. Reasoning about binary topological relations. In Gunther, O. and Schek, H.J. (eds.), *Advances in Spatial Databases*, SSD'91 Proceedings, Springer Verlag. 143–160, 1991.

[Egenhofer and Sharma, 1993] M. J. Egenhofer and Jayant Sharma. Assessing the consistency of complete and incomplete topological information. *Geographical Systems*, 1:47–68, 1993.

[Kirousis and Papadimitriou, 1988] L. Kirousis and C. H. Papadimitriou. The complexity of recognizing polyhedral scenes. *Journal of Computer and Systems Science*, 37(1):14–38, 1988.

[Kratochvíl, 1991] J. Kratochvíl. String graphs II: Recognizing string graphs is NP-hard. *Journal of Comb. Theory, Series B*, 52(1):67–78, 1991.

[Kratochvíl and Matoušek, 1991] J. Kratochvíl and J. Matoušek. String graphs requiring exponential representations. *Journal of Comb. Theory, Series B*, 53(1):1–4, 1991.

[Mackworth and Freuder, 1985] A. Mackworth and E. Freuder. The complexity of some polynomial network consistency algorithms for constraint satisfaction problems. *Artificial Intelligence*, 25:65–74, 1985.

[Papadias and Sellis, 1994] D. Papadias and T. Sellis. The Qualitative Representation of Spatial Knowledge in Two-Dimensional Space. *Very Large Data Bases Journal, Special Issue on Spatial Databases*, 4:479–516, 1994.

[Papadias *et al.*, 1995] D. Papadias, Y. Theodoridis, T. Sellis, and M. Egenhofer. Topological relations in the world of minimum bounding rectangles: a study with R-trees. *Proceedings of the ACM SIGMOD Conference*, San Jose, California, 1995.

[Smith and Park, 1992] Terence R. Smith and Keith K. Park. Algebraic approach to spatial reasoning. *International Journal Geographical Information Systems*, 6:177–192, 1992.