

Local Topological Relationships for Complex Regions

Mark McKenney, Alejandro Pauly, Reasey Praing, and Markus Schneider*

Department of Computer & Information Science & Engineering
University of Florida
Gainesville, FL 32611, USA
{mm7, apauly, rpraing, mschneid}@cise.ufl.edu

Abstract. Topological relationships between spatial objects are important for querying, reasoning, and indexing of data within spatial databases. These relationships are qualitative and respond to questions about the relative positions (e.g., disjointness or containment) of spatial objects. Several models have been proposed that effectively define formal sets of topological relationships between *simple* spatial data types. The generalization of topological relationship models to *complex* spatial data types, which are roughly defined as multi-component versions of their simple counterparts, has raised awareness of the fact that these models only provide a *global* view of topological relationships whereas details of the topological relationships between individual components of the spatial objects involved are often ignored. In this paper, we introduce a fine-grained view on topological relationships between complex regions. Our model focuses on leveraging information about the *local* topological relationships that hold between the components of two spatial objects, thereby providing a *localized* view of the overall global topological relationship.

1 Introduction

The exploration of relationships between spatial objects is an important topic in fields such as robotics, VLSI design, linguistics, CAD, and GIS. Object relationships can be used not only to learn information about the objects involved but also for inferring new non-explicit information as well as creating fast access and indexing structures in spatial databases. Specifically, topological relationships have been the focus of extensive research for a long time. This research includes the design of models of topological relationships between all types of spatial objects as well as related topics like the exploration of topological relationships as a reasoning tool.

Models for topological relationships have predominantly considered *simple* spatial data types. A *simple point* object is defined as a single pair of coordinates, a *simple line* object is given as a non self-intersecting connected curve, and a *simple region* object is represented as an areal object topologically equivalent to a closed disc. A well-known model that defines the topological relationships between simple spatial objects is the *9-intersection* model (9IM). The commonly known set of eight topological relationships originally defined by the 9IM between simple regions includes the relationships *overlap*, *meet*, *inside*, *contains*, *coveredBy*, *covers*, *equal*, and *disjoint*.

* This work was partially supported by the National Science Foundation under grant number NSF-CAREER-IIS-0347574.

More recently, the 9IM has effectively been applied to the latest generation of *complex* spatial data types that, for example, have been propagated by the Open GeoSpatial Consortium. Roughly, a *complex point* is defined by a set of disjoint simple points. A *complex line* is composed of a set of blocks of connected simple lines. A *complex region* is defined as a set of one or more faces, each possibly containing holes. The application of the 9IM to complex spatial data types has raised awareness of the *global* nature of the 9IM. That is, the 9IM considers the interior, exterior, and boundary point sets of the whole objects, and ignores the fact that complex spatial objects are composed of individual and separate components. As a result, *local* topological information regarding the relationship between individual components from each object is lost. Consider a scenario where an oil spill has occurred on a coral reef system. If the oil spill overlaps a section of the reef, then the global relationship would be overlap. However, there may be isolated portions of the reef which are not yet affected by the oil spill, but will be soon if no action is taken. For example, a portion of the reef may meet the edge of the oil spill. The overall topological relationship according to the 9IM will be considered as an *overlap* and the local *meet* relationship will be ignored. If we can identify the local meet interaction, then a barrier can be constructed to save this portion of the reef. This example illustrates the domination of a local overlap over a local meet. But it turns out that this *dominance* problem is not the only way in which local information is lost within the 9IM. For example, an overlap can exist globally when one face of a region *B* is *coveredBy* a face of another region *A* and vice versa a face of *A* is contained in the other face of *B*. Even without the existence of a local overlap, a global overlap appears due to the *composition* of other local relationships.

In this paper, we propose a model for topological relationships between complex regions that is at least as expressive as the 9IM with respect to global information, and at the same time, is able to retain information regarding local relationships between individual components of these regions. We achieve this by defining global topological relationships based on the existence of local topological relationships between the components of the objects involved. This work provides a more fine-grained tool for querying spatial objects and increases the expressive power available to a spatial database user. Furthermore, we provide a notation to specify spatial queries that utilize local topological information.

Section 2 introduces previous related work. Section 3 formally defines the topological relationships between simple regions with holes which are necessary for the formal description of our model. In Section 4, the dominance and composition problems are detailed. Our *locality aware* model of topological relationships between complex regions is presented in Section 5. We consider the expressiveness of our locality aware model compared to the global model in Section 6. In Section 7, we connect our concept with the user by providing the notation that allows our model to be embedded into a common database query language. Finally, Section 8 gives conclusions and future work.

2 Related Work

This section briefly introduces the most important background concepts that are necessary to fully understand and to solve the dominance and composition problems



Fig. 1. (a) Illustrates the interior, boundary, and exterior point sets of a complex region composed of three faces. (b) Illustrates a single face, also denoted as a simple region with holes.

introduced earlier. We concentrate on introducing the definition of complex regions upon which this paper is based. We also describe the 9IM and cover a previous approach that deals with preserving local topological relationships.

We are interested in distinguishing between the interior, boundary and exterior point sets of complex regions, as defined in [1]. Based on this definition, Figure 1(a) illustrates a complex region with its interior, exterior, and boundary.

In order to be able to describe the topological predicates between components of complex regions, we must define their structural components. Complex regions are composed of *faces*. Each face is regarded as a *simple region with holes*, an example of which is illustrated in Figure 1(b). A simple region with holes has a connected interior and (possibly) disconnected exterior and boundary due to the existence of holes. Informally, a simple region with holes is made up of an outer polygon denoting its outer boundary and zero or more hole polygons representing its holes. All holes must be completely contained within the outer polygon and can share a finite number of boundary points with the outer cycle and with other holes. We denote the set of all simple regions with holes as SRH.

Topological relationships between spatial objects can be defined by the 9-intersection model by evaluating the non-emptiness of the intersection between all combinations of the interior ($^{\circ}$), boundary (∂) and exterior ($^{-}$) of the objects involved. A unique 3×3 matrix with Boolean values filled as illustrated in Figure 2 describes the topological relationship between each pair of spatial objects.

Originally defined for simple regions, the 9IM has been extended to handle simple regions with holes [2], complex spatial objects [1], and *composite* regions (i.e., complex regions without holes) [3]. The model in [2] characterizes the topological relationships between two simple regions with holes as the conjunction of topological relationships between their underlying simple regions (the holes are considered simple regions). For two simple regions with holes A and B with n and m holes respectively, a matrix of $(n+1)(m+1)$ elements represents the topological relationship between A and B . This means that under this model, the number of topological relationships between two simple regions with holes is dependent on the number of holes in each region, resulting in an arbitrary number of predicates. To avoid this, we identify the finite set of topological relationships between simple regions with holes based on the 9IM (Section 3) that is independent of the number of holes and that we use as a basis for the rest of this paper.

In [4], the authors introduce a model that preserves local topological relationships between composite regions while still maintaining global information. The approach used to define that model is a precursor to the approach used in this paper, where we

$$\begin{pmatrix} A^\circ \cap B^\circ \neq \emptyset & A^\circ \cap \partial B \neq \emptyset & A^\circ \cap B^- \neq \emptyset \\ \partial A \cap B^\circ \neq \emptyset & \partial A \cap \partial B \neq \emptyset & \partial A \cap B^- \neq \emptyset \\ A^- \cap B^\circ \neq \emptyset & A^- \cap \partial B \neq \emptyset & A^- \cap B^- \neq \emptyset \end{pmatrix}$$

Fig. 2. The 9-intersection matrix for topological relationships

see that the more general problem of modeling local and global topological relationships between complex regions is significantly more complex. Furthermore, in this paper we introduce and define the global topological predicates between simple regions with holes. We also define a local topological relationship specification mechanism for querying spatial objects with local spatial information. This mechanism allows the user to fully exploit locality aware models of topological relationships, and it can be applied to the model presented in this paper as well as the model in [4].

3 Topological Predicates Between Components of Complex Regions

Given that the set of *simple regions with holes* is a subset of the set of complex regions, it follows that the set TP_{srh} of topological relationships that can hold between two simple regions with holes is a subset of the set TP_{cr} of topological relationships that can hold between two complex regions. This is due to the fact that a SRH object represents a specific instance of a complex region; thus, any topological relationship defined for such an object must also be defined for complex regions. This means that we must find the elements (if any) from TP_{cr} that are not members of TP_{srh} . To do so, we impose a new constraint to the relationships denoted by the 9-intersection matrices as originally defined in [1].




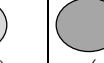


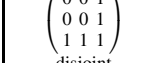
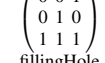
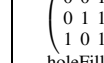
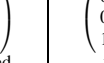
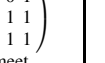
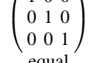
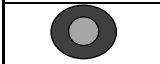
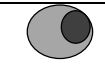
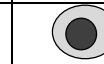
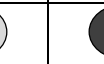

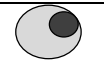
Lemma 1. *Let R and S be two simple regions with holes. If the interior of R intersects both the interior and exterior of S , then the interior of R must also intersect the boundary of S , and vice versa, i.e.,*

$$\begin{aligned} & \forall R, S \in SRH, \\ & (R^\circ \cap S^\circ \neq \emptyset \wedge R^\circ \cap S^- \neq \emptyset \Rightarrow R^\circ \cap \partial S \neq \emptyset) \wedge \\ & (R^\circ \cap S^\circ \neq \emptyset \wedge R^- \cap S^\circ \neq \emptyset \Rightarrow \partial R \cap S^\circ \neq \emptyset) \end{aligned}$$

Proof. It follows from the Jordan-Curve theorem that any point set which intersects both the interior and the exterior of a region but not its boundary must be disconnected. Based on its definition, any SRH must have a connected interior; therefore, it cannot intersect the interior and exterior of another region without intersecting its boundary. \square

This new constraint eliminates 15 of the 33 originally identified topological relationships between complex regions. The remaining 18 relationships represent the members of TP_{srh} . We must now prove that all 18 can actually hold between two SRH objects and that no further constraints are necessary. We follow the method of proof by drawing and show in Table 1 sample configurations of each relationship.

Table 1. The 18 topological predicates between simple regions with holes. One object is shaded dark and the other light as in the *disjoint*, whereas the shared areas have the darkest shade.

 $\begin{pmatrix} 0 & 0 & 1 \\ 0 & 0 & 1 \\ 1 & 1 & 1 \end{pmatrix}$ disjoint	 $\begin{pmatrix} 0 & 0 & 1 \\ 0 & 1 & 0 \\ 1 & 1 & 1 \end{pmatrix}$ fillingHole	 $\begin{pmatrix} 0 & 0 & 1 \\ 0 & 1 & 1 \\ 1 & 0 & 1 \end{pmatrix}$ holeFilled	 $\begin{pmatrix} 0 & 0 & 1 \\ 0 & 1 & 1 \\ 1 & 1 & 1 \end{pmatrix}$ meet	 $\begin{pmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{pmatrix}$ equal	 $\begin{pmatrix} 1 & 0 & 0 \\ 1 & 0 & 0 \\ 1 & 1 & 1 \end{pmatrix}$ inside
 $\begin{pmatrix} 1 & 0 & 0 \\ 1 & 1 & 0 \\ 1 & 0 & 1 \end{pmatrix}$ holeCoveredBy	 $\begin{pmatrix} 1 & 0 & 0 \\ 1 & 1 & 0 \\ 1 & 1 & 1 \end{pmatrix}$ coveredBy	 $\begin{pmatrix} 1 & 1 & 1 \\ 0 & 0 & 1 \\ 0 & 0 & 1 \end{pmatrix}$ contains	 $\begin{pmatrix} 1 & 1 & 1 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{pmatrix}$ holeCovers	 $\begin{pmatrix} 1 & 1 & 1 \\ 0 & 1 & 1 \\ 0 & 0 & 1 \end{pmatrix}$ covers	 $\begin{pmatrix} 1 & 1 & 1 \\ 1 & 0 & 0 \\ 1 & 1 & 1 \end{pmatrix}$ insideOverfill
 $\begin{pmatrix} 1 & 1 & 1 \\ 1 & 0 & 1 \\ 1 & 0 & 1 \end{pmatrix}$ containsOverfill	 $\begin{pmatrix} 1 & 1 & 1 \\ 1 & 0 & 1 \\ 1 & 1 & 1 \end{pmatrix}$ insideContains	 $\begin{pmatrix} 1 & 1 & 1 \\ 1 & 1 & 0 \\ 1 & 0 & 1 \end{pmatrix}$ coversCoveredBy	 $\begin{pmatrix} 1 & 1 & 1 \\ 1 & 1 & 0 \\ 1 & 1 & 1 \end{pmatrix}$ coversOverfill	 $\begin{pmatrix} 1 & 1 & 1 \\ 1 & 1 & 1 \\ 1 & 0 & 1 \end{pmatrix}$ coveredByOverfill	 $\begin{pmatrix} 1 & 1 & 1 \\ 1 & 1 & 1 \\ 1 & 1 & 1 \end{pmatrix}$ overlap

4 Locality of Models for Topological Relationships

In this section, we introduce the problem of locality (or globality) of models for topological predicates that apply to multi-component spatial objects. The degree of locality of a topological predicate model refers to the ability of the model to distinguish between an overall view of a relationship between spatial objects and the relationships between the individual components of those objects. Section 4.1 identifies three degrees of locality and classifies current models of topological relationships according to these degrees. This classification motivates the need to develop a more comprehensive model in terms of locality. This need is illustrated by two main drawbacks found in the 9IM: the *dominance problem*, which we present in Section 4.2, and the *composition problem*, presented in Section 4.3. Once these drawbacks are identified, we sketch the requirements and procedure to derive a locally and globally aware model for topological predicates between complex regions in Section 4.4.

4.1 Views of Topological Predicates

Many models of topological predicates have been proposed in the literature, some of which have been derived from the 9IM model for simple objects; however, there has been little attempt to classify the models in any meaningful way. Here we use an example to motivate a classification of topological predicate models into three *views* of predicates where each view defines a different degree of locality. We provide a single scene of complex region objects (Figure 3) and use it to demonstrate the various views.

The first view of topological predicates we consider is the *global view*. Topological predicates that take into account the global view of a scene equate any given scene to a single topological predicate. For instance, the 9IM for simple regions is a global view model, as is the 9IM for complex regions, which is an extension of the 9IM for simple




Fig. 3. Two complex regions, one darker and one lighter. Their common area is shaded darkest.

regions. The 9IM for complex regions classifies the objects in Figure 3 as satisfying the *overlap* predicate. However, such a global view does not necessarily indicate complete information about a scene. In this case, regions A and B do indeed *overlap*, but a *holeFilled* situation (between A_1 and B_1) also exists, as well as a *disjoint*. While the global view of the 9IM does assign a unique predicate to this scene, it effectively hides other interactions between components of the complex regions.

A second category is the *local view* of topological relationships, which can also be characterized as the *existence view*. This view focuses solely on the topological relationships between the individual components of multi-component objects. Queries that take advantage of the local view can typically be posed in the form of an existence question; for example, a user may wish to ask “does there exist a *meet* interaction between object A and object B ?”. The local view is fundamentally different from the other views in that it does not attempt to assign a unique predicate to a given configuration of objects. Instead, it assigns a unique predicate to the interaction of any two components between spatial objects. In the case of complex regions, these components are simple regions with holes.

Finally, the *hybrid view* combines the information provided by global and local views to allow multiple local interactions in a given scene to be expressed in a unique global predicate. One example of such a view is given in [3], where the authors present a model for topological predicates which describes a scene as a matrix of predicates where each entry in the matrix represents a 9IM predicate between components of each object. Such a view provides a single, although complex, predicate for the entire scene while maintaining the information about all interactions between all faces of the objects. The advantage of the hybrid view is that a global predicate can be assigned to a scene such that local information is not hidden.

4.2 The Dominance Problem

A major drawback to global views of modeling topological relationships is the *dominance problem*. The dominance problem is characterized by a global topological configuration overshadowing the existence of other local topological configurations between a pair of given objects. Figure 3 depicts an arrangement of two complex region objects A and B in the plane. As described above, the regions are globally in an *overlap* configuration due to the overlapping of components A_2 and B_2 . However, if the interaction of each face of object A with object B is examined independently i.e., *locally*, then in addition to the *overlap* configuration, a *holeFilled* configuration is observed between A_1 and B_1 . Because the global configuration for regions is *overlap*, we say that the *overlap* configuration dominates the *holeFilled* configuration. Another dominated relation in the figure is the *disjoint* configuration between A_1 and B_2 .

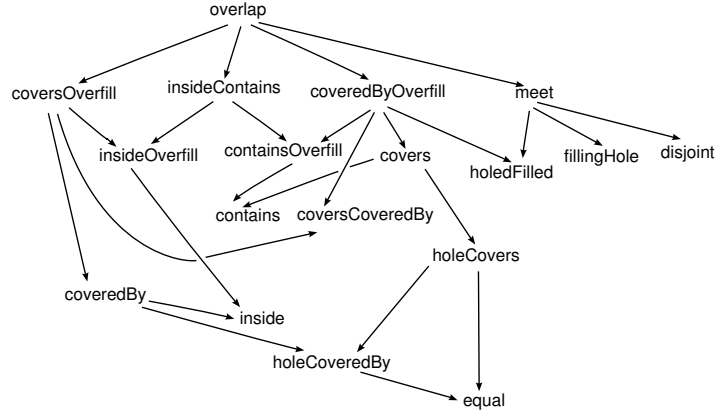


Fig. 4. The dominance hierarchy for topological predicates between simple regions with holes. An arrow $p \rightarrow q$ means that p dominates q .

In order to cope with the dominance problem, it is necessary to determine which topological predicates dominate which other topological predicates. In other words, we must define a dominance order among the topological predicates. We can determine a dominance order based on the 9-intersection matrices of each topological predicate. For a 9-intersection matrix $(p_{ij})_{0 \leq i, j \leq 2}$ representing a topological predicate, we define the *not-empty-entry* set NE_p to be the set of matrix coordinates corresponding to matrix entries that are equal to 1 (*true*). Conversely, we define the *empty-entry* set E_p to be the set of all matrix coordinates corresponding to entries in matrix $(p_{ij})_{0 \leq i, j \leq 2}$ that are equal to 0 (*false*). Formally, we define:

$$NE_p = \{(i, j) | 0 \leq i, j \leq 2 \wedge p_{ij} = 1\}$$

$$E_p = \{(i, j) | 0 \leq i, j \leq 2 \wedge p_{ij} = 0\}$$

Let TP_{srh} be the set of topological predicates between simple regions with holes. For two topological predicates $p, q \in TP_{srh}$, we define the *dominates* operator to return *true* if, and only if, the not-empty-entry set of q is a subset of the not-empty-entry set of p :

$$dominates : TP_{srh} \times TP_{srh} \longrightarrow \mathbb{B}$$

$$dominates(p, q) := \text{if } NE_q \subset NE_p \text{ then } true \text{ else } false$$

Based on the *dominates* operator, for illustration purposes we construct the dominance hierarchy for the elements of the set TP_{srh} of topological predicates between simple regions with holes as shown in Figure 4. It becomes apparent that the dominance problem only affects the topological predicates between multi-component spatial objects; thus, the dominance hierarchy can be determined for the topological predicates between complex points, lines, and regions, as well as any combination of those.

4.3 The Composition Problem

The composition problem arises due to the observation that global topological relationships can be present without actually existing locally. For example in Figure 5, locally




Fig. 5. A scene exposing the composition problem of the 9IM

we distinguish the existence of a *disjoint* between A_1 and B_2 , a *meet* (A_3, B_2), an *equal* (A_3, B_3), a *coveredBy* (A_2, B_2) and a *covers* (A_1, B_1). If we look at the global 9IM describing the topological relationship between A and B , we notice that it represents the *overlap* predicate even though no such *overlap* exists locally. Whereas the dominance problem relates to the fact that the global view hides information about local relationships that do exist, the composition problem hides information about global relationships that do not exist locally.

4.4 Deriving a Hybrid Model for Topological Predicates

We now focus on defining the requirements of a hybrid model for topological predicates between complex regions and developing a procedure to define such a model. The first and obvious requirement is that the model must be able to describe global topological relationships, and at the same time, preserve the necessary information that may be required about the local relationships between components of the objects involved. This is achieved in principle, by defining a model that does not suffer from the dominance or composition problems as described in the previous sections.

The new hybrid model must also be at least as expressive as the 9IM. That is, the new model should be able to distinguish between all (and possibly more) of the relationships that the 9IM between complex regions can distinguish. We define two topological predicate models as equally expressive if there is exactly a one-to-one correspondence between their predicates. Ideally, the new model will have a many-to-one correspondence with the 9IM, yielding one or more predicates for each predicate derived by the 9IM.

The first step in deriving our model is to find a formal characterization of topological predicates between complex regions. This characterization must allow us to describe a topological relationship between two complex regions in such a way that the dominance and composition problems are avoided. Once such a characterization is reached, we must discover all the characterizations that are realizable based on the definition of complex regions. These valid characterizations are now considered topological predicates. Next, we must be able to compare the expressiveness of the 9IM and the new characterizations. This is done in order to ensure that the third requirement set above is fulfilled. The comparison is performed by finding the two-way correspondence between elements in the new characterization and elements (topological predicates) in the global 9IM.

Once we have ensured that the new model is as expressive as the 9IM globally, we must provide a mechanism that will allow the user to exploit the local information retained by the new characterization. This ensures that the hybrid model can be used both from a global, and from a local perspective.

5 A Model for Preserving Local Interactions

In this section, we present a novel model of *localized topological relationships* for multi-component spatial objects that does not suffer from the dominance and composition problems. The *localized topological predicate* (LTP) model is a hybrid view of topological relationships in that it utilizes local information to provide a global relationship between two objects. We begin by exploring the characterization of LTP, and then discovering all LTPs that are possible between complex regions. Finally we compare this new model with the 9IM between complex regions and attempt to reach a completely hybrid (local and global) model based on this comparison.

5.1 Characterization of Localized Topological Predicates

In order to discover the local topological relationships between complex regions, we define *existence topological predicates* (ETP). Let cr be the set of all complex regions, and ET_{cr} be the set of all existence topological predicates between complex regions. For $p \in TP_{srh}$, the ETP $p_e \in ET_{cr}$ is defined as a function $p_e : cr \times cr \rightarrow boolean$. For any $A, B \in cr$:

$$p_e(A, B) = \begin{cases} true, & \exists A_i \in A, \exists B_j \in B : p(A_i, B_j) = true \\ false, & otherwise \end{cases}$$

Based on this definition, for complex regions we identify a total of 18 existence topological predicates that exactly correspond to the 18 topological predicates between simple regions with holes, i.e., $ET_{cr} = \{disjoint_e, fillingHole_e, holeFilled_e, meet_e, equal_e, inside_e, holeCoveredBy_e, coveredBy_e, contains_e, holeCovers_e, covers_e, insideOverfill_e, containsOverfill_e, insideContains_e, coversCoveredBy_e, coversOverfill_e, coveredByOverfill_e, overlap_e\}$.

For complex regions A and B we define the localized topological predicate that describes their relationship as a conjunctive boolean expression with exactly eighteen clauses, each with a single element that corresponds to a predicate from ET_{cr} or its negation. That is, an LTP characterizes the topological predicate between two complex regions by asserting which ETPs hold and which do not hold between the complex regions. Let $E(A, B)$ and $F(A, B)$ be the sets of ETPs that yield *true* and *false* respectively, for A and B , i.e.

$$\begin{aligned} E(A, B) &= \{e \in ET_{cr} \mid e(A, B) = true\} \\ F(A, B) &= \{f \in ET_{cr} \mid f(A, B) = false\} \end{aligned}$$

We define $l \in LT_{cr}$ (where LT_{cr} is the set of LTPs between complex regions) as:

$$l(A, B) = \left(\bigwedge_{e \in E(A, B)} e(A, B) \right) \wedge \left(\bigwedge_{f \in F(A, B)} \neg f(A, B) \right)$$

For notation purposes, we represent each LTP as a 18-bit vector where each bit corresponds to an existence predicate in the order shown in Table 1. Based on this definition, we denote the set of all valid 18-bit vectors as $LTBV_{cr}$. The element from this set that represents the scenario in Figure 5 is $[1, 0, 0, 1, 1, 0, 0, 1, 0, 0, 1, 0, 0, 0, 0, 0, 0, 0]$,

because the local topological relationships that exist between components in the scene are *disjoint*, *meet*, *equal*, *coveredBy*, and *covers*.

5.2 Identifying the Valid Characterizations

In order to identify the elements that belong to LT_{cr} , we first consider the complete set of 18-bit vectors that can exist (i.e., 2^{18} such vectors). By considering the semantics of the bits, we can see that not all combinations are possible. A trivial example is represented by the *all-zeros* TBV that we can determine to be topologically invalid due to the fact that every pair of (non-empty) simple regions with holes must satisfy at least one of the eighteen topological predicates. We follow a procedure of *constraint* and *validation* in order to successfully identify the complete set of valid TBVs for complex regions. The first step of the procedure is to eliminate, by way of constraint rules, TBVs that are topologically invalid. The second step entails validating the remaining TBVs. If all remaining TBVs are successfully validated, then we have reached a complete set, otherwise we must identify a new constraint rule. This sequence is repeated until the validation is complete.

The constraint rules for the first step are expressed by way of lemmas that provide formal proof that certain situations within bit vectors cannot concurrently occur. For example, take the complex regions A and B such that a component (simple region with hole) A_i of A and a component B_j of B are *equal*. Besides this local relationship, there exists a local *overlap* relationship between some other component A_k of A and B_l of B . For such a configuration, we know that so far the *equal* and *overlap* bits are set to 1. But we can also determine that either the *disjoint* or the *meet* bits must be set to account for the relationships between A_i and B_l and between A_k and B_j . It must be true that $(A_i \neq A_k) \wedge (B_j \neq B_l)$, otherwise two components from the same complex region object would have intersecting interiors which are disallowed by the definition of complex regions. Formally,

Lemma 2. *Let $A, B \in SRH$, let A_i, A_k, B_j , and B_l be faces of A and B , respectively, and let $q, r \in TP_{srh}$:*

$$\begin{aligned} \text{equal}(A_i, B_j) \wedge q(A_k, B_l) &\Rightarrow r(A_i, B_l) \text{ where} \\ &(q(A_k, B_l) \Rightarrow A_k^\circ \cap B_l^\circ \neq \emptyset) \wedge r \in \{\text{disjoint}, \text{meet}\} \end{aligned}$$

$$\begin{aligned} \text{Proof. } \text{equal}(A_i, B_j) &\Rightarrow A_i^\circ = B_j^\circ \\ q(A_k, B_l) &\Rightarrow A_k^\circ \cap B_l^\circ \neq \emptyset \\ A_i^\circ = B_j^\circ \wedge A_k^\circ \cap B_l^\circ \neq \emptyset &\Rightarrow A_i^\circ \cap B_l^\circ = \emptyset \\ &\Rightarrow \text{disjoint}(A_i, B_l) \vee \text{meet}(A_i, B_l) \quad \square \end{aligned}$$

We have identified a total of 17 lemmas, including Lemma 2 above and the trivial case of an all zeros bit vector (Lemma 0). All lemmas other than the trivial all zero case are of the form $p(A_i, B_j) \wedge q(A_k, B_l) \Rightarrow r(A_i, B_l)$. Table 2 includes the definitions of p , q , and r for the remaining 15 lemmas. The proofs for all these follow a structure similar to the proof of Lemma 2; thus, we choose to omit such proofs due to space constraints.

The completeness of these lemmas is validated by a mechanism designed to ensure that each 18-bit vector that is not covered by the constraints, is realizable and valid. The mechanism is based on concepts of topological reasoning that are applied through a theory of *Binary Constraint Networks* (BCN). In the mechanism we are able to reason about topological relationships by computing the composition operation [5]. The

Table 2. The remaining lemmas acting as constraints for the invalidation of topological bit vectors

Lemma	p	q	r
3	$p \in \{\textit{inside}\}$	$q(A_k, B_l) \Rightarrow \overline{A_k} \supseteq \overline{B_l}$	$r \in \{\textit{disjoint}\}$
4	$p \in \{\textit{contains}\}$	$q(A_k, B_l) \Rightarrow \overline{A_k} \subseteq \overline{B_l}$	$r \in \{\textit{disjoint}\}$
5	$p \in \{\textit{coveredBy}\}$	$q(A_k, B_l) \Rightarrow A_k^\circ \supseteq \partial B_l$	$r \in \{\textit{disjoint}, \textit{meet}\}$
6	$p \in \{\textit{covers}\}$	$q(A_k, B_l) \Rightarrow \partial A_k \subseteq B_l^\circ$	$r \in \{\textit{disjoint}, \textit{meet}\}$
7	$p \in \{\textit{overlap}\}$	$q(A_k, B_l) \Rightarrow \partial A_k \cap \partial B_l = \emptyset \wedge A_k^- \cap \partial B_l = \emptyset$	$r \in \{\textit{disjoint}, \textit{meet}\}$
8	$p \in \{\textit{fillingHole}\}$	$q(A_k, B_l) \Rightarrow A_k^- \cap \partial B_l = \emptyset$	$r \in \{\textit{disjoint}, \textit{meet}\}$
9	$p \in \{\textit{holeFilled}\}$	$q(A_k, B_l) \Rightarrow \partial A_k \cap B_l^- = \emptyset$	$r \in \{\textit{disjoint}, \textit{meet}\}$
10	$p \in \{\textit{holeCoveredBy}\}$	$q(A_k, B_l) \Rightarrow \partial A_k \cap \partial B_l \neq \emptyset \vee \partial A_k \cap B_l^\circ \neq \emptyset$	$r \in \{\textit{disjoint}, \textit{meet}\}$
11	$p \in \{\textit{holeCovers}\}$	$q(A_k, B_l) \Rightarrow \partial A_k \cap \partial B_l \neq \emptyset \vee A_k^\circ \cap \partial B_l \neq \emptyset$	$r \in \{\textit{disjoint}, \textit{meet}\}$
12	$p \in \{\textit{insideOverfill}\}$	$q(A_k, B_l) \Rightarrow A_k^- \cap \partial B_l = \emptyset \wedge A_k^\circ \cap B_l^- \neq \emptyset$	$r \in \{\textit{disjoint}\}$
13	$p \in \{\textit{containsOverfill}\}$	$q(A_k, B_l) \Rightarrow \partial A_k \cap B_l^- = \emptyset \wedge A_k^\circ \cap B_l^- \neq \emptyset$	$r \in \{\textit{disjoint}\}$
14	$p \in \{\textit{insideContains}\}$	$q(A_k, B_l) \Rightarrow A_k^- \cap \partial B_l = \emptyset \wedge A_k^- \cap \partial B_l = \emptyset$	$r \in \{\textit{disjoint}, \textit{meet}\}$
15	$p \in \{\textit{coversCoveredBy}\}$	$q(A_k, B_l) \Rightarrow A_k^- \cap B_l^\circ \neq \emptyset \wedge A_k^\circ \cap B_l^- \neq \emptyset$	$r \in \{\textit{disjoint}, \textit{meet}\}$
16	$p \in \{\textit{coversOverfill}\}$	$q(A_k, B_l) \Rightarrow A_k^- \cap \partial B_l = \emptyset \wedge A_k^- \cap B_l^\circ \neq \emptyset$	$r \in \{\textit{disjoint}, \textit{meet}\}$
17	$p \in \{\textit{coveredByOverfill}\}$	$q(A_k, B_l) \Rightarrow \partial A_k \cap B_l^- = \emptyset \wedge A_k^\circ \cap B_l^- \neq \emptyset$	$r \in \{\textit{disjoint}, \textit{meet}\}$

composition of two relationships p and q between A and B , and between B and C respectively, is able to at least partially derive the relationship r between A and C . A common way to denote this operation is $p(A, B); q(B, C) \Rightarrow r(A, C)$.

We model each of the 18-bit vectors as a *Binary Spatial Constraint Network* (BSCN). A BSCN is considered valid if it is *path-consistent* [6]. Path-consistency is computed by using the composition of topological relationships between simple regions with holes. For example take a BSCN M where M_{ij} refers to the constraints (relationships) between variables (objects) i and j . We say that M is path-consistent if it holds that $M_{ij} \subseteq M_{ik}; M_{kj}$ for every i, j , and k variable of M .

To model a bit vector as a BSCN we consider a scene made up of two complex regions A and B with n and m simple region with holes components respectively such that $A = \{A_0, A_1, \dots, A_n\}$, and $B = \{B_0, B_1, \dots, B_m\}$. A BSCN is a graph constructed by considering each component as a variable (vertex of graph) and then assigning constraints between each variable as the topological relationships represented in the edges of the graph. The edges between two components of the same object are implicitly defined by the definition of complex regions, whereas the rest of edges are defined by the bit vector that is validated. For each bit vector we construct several scenarios with different numbers of components, and for each scenario we iterate over all possible assignments of predicates matching the bit vector. This is repeated for each bit vector until it is either validated or a constraint can be proven that invalidates the vector.

As seen in Table 2, for complex regions we have identified the 17 constraints which eliminate all the invalid bit vectors. The final set LT_{cr} is composed of 137209 18-bit vectors.




Fig. 6. Two spatial configurations with the same topological bit vector but a different 9-intersection matrix

5.3 Comparing the Models

Although the $LTBV_{cr}$ model of localized topological relationships effectively exposes the local interactions among components of complex regions, it is not yet clear how expressive the model is in comparison to the 9IM. We determine the expressiveness of the LTP model for complex regions with respect to the 9IM based on the following: if every spatial configuration that is representable and distinguishable by the 9IM can also be represented and distinguished by the LTP model, then the LTP model is at least as *globally* expressive as the 9IM. Otherwise, we say that the LTP model is *globally* less expressive than the 9IM.

It turns out that a number of spatial configurations are distinguishable by the 9-intersection model (i.e., they have different matrix representations) but not by the LTP model. For example, both scenes in Figure 6 are represented by different 9-intersection matrices but by the single bit vector $[0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 1]$, i.e., all pairs of components from the respective objects *overlap*. Thus, we conclude that the LTP model is *globally* less expressive than the global 9IM.

Conversely, we compare the *local* expressiveness of the models to determine whether every spatial configuration that is representable and distinguishable by the LTP model can also be represented and distinguished by the 9IM. Figures 7(a), and 7(b) illustrate two spatial configurations which are distinguished by two different bit vectors $[1, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 1]$, and $[0, 0, 0, 1, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 1]$ respectively. However, both scenes are represented only by a single 9-intersection matrix. Thus, we conclude that the 9IM is *locally* less expressive than the LTP model.

Due to the differences in expressiveness of the LTP model and the 9IM at both the global and local levels we cannot draw any conclusion as to the relative general expressiveness of the models. This is due to the fact that the models are fundamentally incomparable. Our goal is to produce a hybrid model of topological predicates that is at least as globally expressive as the 9IM. We observe that the scenes in Figure 6 cannot be differentiated due to the fact that a global matrix can include global information that cannot be determined from examining purely local information. For example, the bit vector for the scenes in Figure 6 merely indicates that there exists a local *overlap* between a component of complex region A and a component of complex region B , whereas both scenes are distinguished by two different global matrices because of the difference in the existence of intersection between the boundary of A and the exterior of B . The non-existence of this intersection in Figure 6a occurs because another local *overlap* exists such that the boundary of A is now part of the closure of B . To determine such global information from a bit vector, more than simply local information must be included in the LTP model.




Fig. 7. Two spatial configurations with the same 9-intersection matrix but a different TBV

6 Maintaining Global Information

Since we have shown the incomparability of the expressiveness between the LTP and the 9IM, we must determine a different, hybrid view model that is at least as globally expressive as the 9IM model. As described in Section 5.3, some global information cannot be expressed by the LTP model, which keeps purely local information. In order to handle this issue, the necessary global information must also be represented in the model. In this section, we present a *hybrid topological predicate* (HTP) model that combines both local and global information into a single predicate. We refer to the set of all HTPs between complex regions as HT_{cr} . We begin by determining which global information must be included into the HTP model, and then define a new bit vector for HTPs and use it to derive a complete set of valid HTPs between complex regions. Finally, we compare the 9IM and the HTP model.

6.1 Hybrid Characterization of Topological Predicates

In order to determine which global information is needed in the HTP model, we examine the effects of local interactions on global interactions, and vice versa. To do this, we identify the collections of matrix entries in the global matrix that must be 1 (non-empty intersection) in order for a given local interaction to be possible. In other words, if an existence predicate yields *true*, we must define how it is expressed in the global matrix. For example, if a *meet* interaction occurs between any pair of components from two complex regions, then we know that in the global matrix representing the scene, the boundaries of the objects must intersect. In general, if any local interaction exists between two complex regions such that the interiors intersect, the boundaries intersect, or the interior intersects the boundary, then the corresponding matrix entry in the global 9-intersection matrix will be 1. This is because if any of these interactions occur locally, the existence of other local interactions cannot cause them to not be reflected at the global level. In contrast, any local interactions involving the exterior of either region may not be reflected in the global matrix because other local interactions can overwhelm them. For example, in Figure 6a, the interiors and boundaries of A_1 and B_1 intersect, and adding other local interactions cannot cause them to no longer intersect; thus, they will intersect in the global matrix. Locally, A_1 and B_1 *overlap*, thus their boundaries and interiors intersect with each other's exteriors. However, the inclusion of B_2 in the scene causes the interaction involving the boundary of A_1 with the exterior of B_1 to not be reflected in the global matrix. This is because globally, the exterior of B_1 is not considered independently, instead the exterior of B is considered as a whole.

We determine that the global information that is required to distinguish the scenes in Figure 6 from each other involves only the interactions of the exteriors of the objects. Based on [1], we know that the intersection between exteriors is always non-empty. Based on the observations detailed above, we must be able to represent the other four exterior interactions (i.e., intersections between interior and exterior, boundary and exterior, and their converses) in order to achieve a *one-to-many* relationship from the 9IM to the HTP model so that all cases that the 9IM can differentiate can also be uniquely identified by the HTP model.

To support the HTP model, we extend the topological bit vector into a new *hybridized topological bit vector* whose bit entries represent both local information and global information. The local information bits represent the eighteen ETPs between complex regions described earlier, whereas the global information bits represent the global exterior interactions. We denote the four meaningful global interactions involving the exterior of either object as *global interaction predicates* (GIPs) and define them as: $ie = (A^\circ \cap B^- \neq \emptyset)$, $be = (\partial A \cap B^- \neq \emptyset)$, $ei = (A^- \cap B^\circ \neq \emptyset)$, and $eb = (A^- \cap \partial B \neq \emptyset)$. GIPs are expressed by four global information bits in the hybridized topological bit vector. Thus, each element of HT_{cr} is represented by a 22-bit vector where the first 18 bits represent the local interactions, and the last four represent the GIP in the order we have just presented them. A bit in the vector is set to 1 if its corresponding ETP or GIP yields *true* and 0 otherwise. The set of all hybridized topological bit vectors is denoted $HTBV_{cr}$. Now, we can distinguish between the scenes in Figure 6: the bit vectors corresponding to Figure 6a and 6b are $[0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 1, 0, 0, 1, 1]$, and $[0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 1, 1, 1, 1, 1]$ respectively.

6.2 Ensuring Expressive Power

The final step in the definition of hybridized topological predicates is the identification of the set $HTBV_{cr}$ that represents the elements in HT_{cr} . We make the assertion that for the elements in $HTBV_{cr}$, the first 18 bits (local information) of all 22-bit vectors must match one of the elements in $LTBV_{cr}$ identified following the method described in Section 5.2. This assertion is based on the fact that the information contained in those bits must be identical to an element in $LTBV_{cr}$ or else it describes a topologically invalid scene. The addition of GIPs only reflects global information and thus does not affect the validity of the topological configuration described by the first 18 bits.

To identify the set of valid 22-bit vectors we provide a two stage method for determining the relationship between the set of 9-intersection matrices and the set of 22-bit vectors: (1) determine the set of 9-intersection matrices that correspond to each 18-bit configuration, and (2) create a 22-bit vector for each corresponding 18-bit vector and 9-intersection matrix by taking the four bits ie , be , ei , and eb of the 9-intersection matrix and extending the 18-bit vector into a 22-bit vector using these GIPs.

We achieve stage one using the matrix templates shown in Table 3. For each ETP we derive a set of template matrices so that each template shows how the existence of a local interaction can be expressed in a global 9-intersection matrix. The values of the templates are based on two observations: first, interactions between the exterior of one component of a complex region and the interior or boundary of a component of another complex region may not be expressed globally due to the existence of other local

Table 3. Templates for mapping 9IM to HTP. Hyphens indicate “don’t care” values.

$\begin{pmatrix} - & - & - \\ - & - & - \\ - & - & 1 \end{pmatrix}$ <i>disjoint_e</i>	$\begin{pmatrix} - & - & 1 \\ - & 1 & - \\ 1 & - & 1 \end{pmatrix}$ <i>fillingHole_e</i>	$\begin{pmatrix} - & - & 1 \\ - & 1 & - \\ 1 & - & 1 \end{pmatrix}$ <i>holeFilled_e</i>	$\begin{pmatrix} - & - & - \\ - & 1 & - \\ - & - & 1 \end{pmatrix}$ <i>meet_e</i>	$\begin{pmatrix} 1 & - & - \\ - & 1 & - \\ - & - & 1 \end{pmatrix}$ <i>equal_e</i>	$\begin{pmatrix} 1 & - & - \\ 1 & - & - \\ 1 & - & 1 \end{pmatrix}$ <i>inside_e</i>
$\begin{pmatrix} 1 & - & - \\ 1 & 1 & - \\ 1 & - & 1 \end{pmatrix}$ <i>holeCoveredBy_e</i>	$\begin{pmatrix} 1 & - & - \\ 1 & 1 & - \\ 1 & - & 1 \end{pmatrix}$ <i>coveredBy_e</i>	$\begin{pmatrix} 1 & 1 & 1 \\ - & - & - \\ - & - & 1 \end{pmatrix}$ <i>contains_e</i>	$\begin{pmatrix} 1 & 1 & 1 \\ - & 1 & - \\ - & - & 1 \end{pmatrix}$ <i>holeCovers_e</i>	$\begin{pmatrix} 1 & 1 & 1 \\ - & 1 & - \\ - & - & 1 \end{pmatrix}$ <i>covers_e</i>	$\begin{pmatrix} 1 & 1 & 1 \\ 1 & - & - \\ 1 & - & 1 \end{pmatrix}$ <i>insideOverfill_e</i>
$\begin{pmatrix} 1 & 1 & 1 \\ 1 & - & - \\ 1 & - & 1 \end{pmatrix}$ <i>containsOverfill_e</i>	$\begin{pmatrix} 1 & 1 & 1 \\ 1 & - & - \\ 1 & - & 1 \end{pmatrix}$ <i>insideContains_e</i>	$\begin{pmatrix} 1 & 1 & 1 \\ 1 & 1 & - \\ 1 & - & 1 \end{pmatrix}$ <i>coversCoveredBy_e</i>	$\begin{pmatrix} 1 & 1 & 1 \\ 1 & 1 & - \\ 1 & - & 1 \end{pmatrix}$ <i>coversOverfill_e</i>	$\begin{pmatrix} 1 & 1 & 1 \\ 1 & 1 & - \\ 1 & - & 1 \end{pmatrix}$ <i>coveredByOverfill_e</i>	$\begin{pmatrix} 1 & 1 & 1 \\ 1 & 1 & - \\ 1 & - & 1 \end{pmatrix}$ <i>overlap_e</i>

interactions. This argument is exposed in the previous section as the reason for the GIPs. As an example of this observation, in Figure 6a, the local interaction between the exterior of B_1 and the boundary of A_1 is not expressed globally because of the existence of B_2 . Therefore, for each template all given values of 1 signify that the value must be set in the global matrix if the existence predicate bit is set in the bit vector. Values for which no global assumption can be made based on the local predicate are denoted in the template with a dash.

The second observation used to derive the templates is that the existence of a local interaction involving the interior or boundary of a component of one complex region and the interior or boundary of a component of another complex region asserts that the corresponding interaction will be reflected in the global matrix describing the scene. Thus, if an ETP is *true* for a scene and it involves one of these four interactions, then all of its corresponding matrix templates will also have a 1 in the corresponding matrix entry. The first observation leads us to denote template values that may *not* occur even if they occur in a local matrix from the same scene, whereas the second observation leads us to denote template values that must occur in the global matrix if they occur in any local matrix. Once the templates are computed, we compute all possible permutations of template matrices where dashes in the entries involving the exterior of a region are replaced with values of 1 such that the *non-empty-entry* set of the template is a subset of the *non-empty-entry* set of its corresponding 9-intersection matrix. These permutations show all possible ways a local interaction may be reflected in a global matrix.

Using the matrix templates, we construct a correspondence between global matrices and 18-bit configurations as follows. First, we consider the set K_m of all matrix templates whose *non-empty-entry* set is a subset of the *non-empty-entry* set of a given 9-intersection matrix m (the *non-empty-entry* set of a template matrix consists of all entries that are set to 1). Then we consider all elements of the powerset of K_m where their boolean sum¹ is equivalent² to m . For each of these elements, we create an 18-bit configuration in which a corresponding ETP has a value of 1 if any of its templates exist in the combination. If the resulting 18-bit configuration does not match one of the identified

¹ If an entry is equal to 1 in any matrix in the set, it is equal to 1 in the boolean sum of the matrices of that set.

² We consider two matrices to be equivalent if their *non-empty-entry* sets are equal.

topologically valid 18-bit vectors, it is discarded. The non-discarded configurations describe the same topological scene as the given 9-intersection matrix m .

In the second stage, we extend each 18-bit configuration to a 22-bit configuration by assigning the combination of GIPs from each 9-intersection matrix to the global information bits of the corresponding 22-bit configurations. This results in a set of 22-bit vectors that represent the elements in $HTBV_{cr}$. We know this set to be complete based on two arguments: first, the 18-bit configurations are proved to be complete based on the constraint and validation mechanism. Second, the values of the global information bits in a 22-bit vector are directly derived from the 9-intersection matrices that map to the corresponding 18-bit configuration. Any other combination of values for the global information bits would force a mapping into a global matrix that has been proven topologically invalid in Section 3.

Besides identifying the elements of HT_{cr} , the method described above also identifies a *one-to-many* relationship between a 9-intersection matrix and elements of $HTBV_{cr}$. For this reason we determine that the hybridized topological predicate model can distinguish more topological scenes than the 9IM, and is therefore more *globally* expressive than the 9IM while maintaining local information.

7 Querying with Local Information

In this section, we introduce a proposal for making localized topological predicates available to a user intending to perform queries involving complex regions. We focus on leveraging features of a common database query language, SQL, to bring the LTP concept closer to the user. Section 7.1 introduces the notation that allows users to pre-define locality aware predicates that can be used in queries. In Section 7.2 we present sample queries that can make use of these predefined predicates but that can also be specified by embedding their definition within the query.

7.1 Constructing Locality Aware Topological Predicates

For the design of a syntactic construction for defining locality aware topological predicates, we focus on what the user may know and may want to express in her queries. The user needs not be aware of the underlying topological bit vector. Instead, the user simply wants to pose a query based on local topological relationships. For example, the user may want to retrieve all pairs of regions that have local situations of *disjoint* and *overlap* but where no *meet* situations occur. Thus, the user is specifying what *must* occur (*disjoint*, *overlap*), and what *must not* occur (*meet*). The rest of the local predicates

<pre>CREATE LTP ON <type1>, <type2> AS <name> ([MUST ALL (<predlist1>)], [MUST NOT ALL (<predlist2>)], [MUST ONE (<predlist3>)], [MUST NOT ONE (<predlist4>)]])</pre>	<pre>CREATE LTP ON region, region AS surrounded(MUST ONE (equal, inside, contains), MUST NOT ALL (overlap))</pre>	<pre>CREATE LTP ON region, region AS meeting(MUST ONE (holeFilled, fillingHole, meet))</pre>
(a)	(b)	(c)

Fig. 8. SQL syntax and examples of DDL for predefining locality aware topological predicates




Fig. 9. Sample illustration and sample query using locality aware topological predicates

would be assessed as optional, meaning that they *may* (or may not) occur. In addition to specifying local existence predicates, the user should be able to specify any of the global interaction predicates that must, must not, and may occur. In implementation, the user can make use of the construct `CREATE LTP` which is defined as part of the query language's *data definition language* (DDL) and whose syntax is shown in Figure 8(a).

In the syntactic definition in Figure 8(a), `type1` and `type2` refer to the spatial data types upon which the predicates operate; these are currently only (complex) regions. The parameter name refers to a user defined label for the predicate that is being created. There are four optional clauses in the construct, divided into two groups: the ALL clauses and the ONE clauses. Each group has two options, the MUST and the MUST NOT. For the MUST clauses the user specifies a list of local predicates and GIPs that must appear to make the predefined predicate *true*. In the case of the MUST ALL clause, all local predicates and GIPs included in the clause must appear. For the MUST ONE clause, at least one of the predicates must appear. Similarly, the MUST NOT clauses allows the user to specify local predicates and GIPs that must not appear for the predicate to result in *true*. It also has two flavors: MUST NOT ALL, and MUST NOT ONE. All local predicates and GIPs that are not included in any clause are implicitly considered to be optional, thus they may or may not appear without changing the outcome of the predicate.

7.2 Local Topological Queries

Having defined a mechanism that allows the user to predefine locality aware topological predicates on complex regions, we move on to show two examples of how this predefined predicates can be used. We use Figure 9 to illustrate both sample queries. For the first example, assume that the elements in Figure 9(a) are complex regions that represent the roaming areas of two distinct species of animals. Although the areas of both species globally overlap, the user is interested in finding animals whose roaming areas somehow meet, as defined by the LTP *meeting* predefined in the previous section. Such a query would use the predefined predicate in the WHERE clause as shown in Figure 9(b).

From this example we notice how predefined LTPs can be used just as existing Boolean predicates. A system implementing LTPs should have system predefined predicates for all local topological predicates operating on the components of supported types. This means that the user may use locality aware existence predicates such as *fillingHole* and *containsOverflow* for a system supporting complex regions and their component simple regions with holes.

In lieu of predefining LTPs, users can also define them *inline* within the WHERE clause. The syntax definition of the sub-clause that is necessary for the inline definition is shown in Figure 10(a). In Figure 10(b) we show how this inline mechanism can be

<pre>WHERE LTP on <attr1>, <attr2> ([MUST ALL (<predlist1>)], [MUST NOT ALL (<predlist2>)], [MUST ONE (<predlist3>)], [MUST NOT ONE (<predlist4>)])</pre>	<pre>SELECT a.name, b.name FROM animals a, animals b WHERE a.avgsize<50 AND b.avgsize>50 AND b.diet=CARNIVORE AND LTP on a.roam, b.roam (MUST ONE (fillingHole, inside, coveredBy))</pre>
(a)	(b)

Fig. 10. Inline use of locality aware predicates within SQL queries: (a) syntax, (b) query

exploited to retrieve all animals and their predators whose average weight is less than 50 lbs. and whose roaming areas are surrounded by the roaming areas of larger carnivores.

8 Conclusion

In this paper, we have defined the concept of local topological predicates between complex regions. We have provided details for the identification of valid LTPs, and shown how these LTPs can be exploited by users through querying. The query language extensions described allow the users to define their own locality aware topological predicates that either can test for the existence of certain local interactions, or can be used to group LTPs together based their common attributes. The result of this work is a new model of topological predicates that is more expressive than the global 9IM and does not suffer from the dominance and composition problems.

Future work on this concept includes the actual implementation of LTPs within a spatial database and its application in real world scenarios. Another interesting topic of related research is the application of the LTP concept as part of indexing mechanisms for complex spatial objects.

References

1. Schneider, M., Behr, T.: Topological Relationships between Complex Spatial Objects. *ACM Trans. on Database Systems* 31, 39–81 (2006)
2. Egenhofer, M., Clementini, E., Di Felice, P.: Topological Relations between Regions with Holes. *Int. Journal of Geographical Information Systems* 8, 128–142 (1994)
3. Clementini, E., Di Felice, P., Califano, G.: Composite Regions in Topological Queries. *Information Systems* 20, 579–594 (1995)
4. McKenney, M., Pauly, A., Praing, R., Schneider, M.: Preserving Local Topological Relationships. In: *ACM Symp. on Geographic Information Systems*, pp. 123–130. ACM Press, New York (2006)
5. Egenhofer, M.J.: Deriving the Composition of Binary Topological Relations. *Journal of Visual Languages and Computing* 2, 133–149 (1994)
6. Ladkin, P., Maddux, R.: On Binary Constraint Problems. *Journal of the Association for Computing Machinery* 41, 435–469 (1994)