

Consistent Query Answering under Spatial Semantic Constraints

M. Andrea Rodríguez
Universidad de Concepción, Chile
andrea@udec.cl

Leopoldo Bertossi*
Carleton University, Canada
bertossi@scs.carleton.ca

Mónica Caniupán
Universidad del Bío-Bío, Chile
mcaniupa@ubiobio.cl

Abstract

Consistent query answering is an inconsistency tolerant approach to obtaining semantically correct answers from a database that may be inconsistent with respect to its integrity constraints. In this work we formalize the notion of consistent query answer for spatial databases and spatial semantic integrity constraints. In order to do this, we first characterize conflicting spatial data, and next, we define admissible instances that restore consistency while staying close to the original instance. In this way we obtain a repair semantics, which is used as an instrumental concept to define and possibly derive consistent query answers. We then concentrate on a class of spatial denial constraints and spatial queries for which there exists an efficient strategy to compute consistent query answers. This study applies inconsistency tolerance in spatial databases, rising research issues that shift the goal from the consistency of a spatial database to the consistency of query answering.

1 Introduction

Consistency in database systems is defined as the satisfaction by a database instance of a set of integrity constraints (ICs) that restricts the admissible database states. Although consistency is a desirable and usually enforced property of databases, it is not uncommon to find inconsistent spatial databases due to data integration, unforced integrity constraints, legacy data, or time lag updates. In the presence of inconsistencies, there are alternative courses of action: (a) ignore inconsistencies, (b) restore consistency via updates on the database, or (c) accept inconsistencies, without changing the database, but compute the “consistent or correct” answers to queries. For many reasons, the first two alternatives may not be appropriate [6], specially in the case of virtual data integration [5], where centralized and global changes to the data sources are not allowed. The latter alternative has been investigated in the relational case [4, 10]. In this paper we explore this approach in the spatial domain, i.e., for spatial databases and with respect to spatial semantic integrity constraints (SICs).

Extracting consistent data from inconsistent databases could be qualified as an “inconsistency tolerant” approach to querying databases. A piece of data will be part of a consistent answer if it is not logically related to the inconsistencies in the database with respect to its set of ICs. We introduce this idea using an informal and simple example.

*Faculty Fellow of the IBM Center for Advanced Studies. Also affiliated to Universidad de Concepción, Chile.

Example 1 Consider a database instance with a relation *LandP*, denoting land parcels, with a thematic attribute (*idl*), and a spatial attribute, *geometry*, of data type *polygon*. An IC stating that geometries of two different land parcels must be disjoint or just touch, i.e., they cannot internally intersect, is expected to be satisfied. However, the instance in Figure 1 does not satisfy this IC and therefore it is inconsistent: the land parcels with ids *idl*₂ and *idl*₃ overlap. Notice that these geometries are partially in conflict and what is not in conflict can be considered as consistent data.



Figure 1: An inconsistent spatial database

Suppose that a query requests all land parcels whose geometries intersect with a query window, which represents the spatial region shown in Figure 1 as a rectangle with dashed borders. Although the database instance is inconsistent, we can still obtain useful and meaningful answers. In this case, only the intersection of *g*₂ and *g*₃ is in conflict, but the rest of both geometries can be considered consistent and should be part of any “database repair” if we decide to restore consistency by means of minimal geometric changes. Thus, since the non-conflicting parts of geometries *g*₂ and *g*₃ intersect the query window, we would expect an answer including land parcels with identities *idl*₁, *idl*₂ and *idl*₃. □

If we just concentrate on (in)consistency issues in databases (leaving aside consistent query answering for a moment), we can see that, in contrast to (in)consistency handling in relational databases, that has been largely investigated, not much research of this kind has been done for spatial databases. In particular, there is not much work around the formalization of semantic spatial ICs, satisfaction of ICs, and checking and maintenance of ICs in the spatial domain. However, some papers address the specification of some kinds of integrity constraints [8, 20], and checking topological consistency at multiple representations and for data integration [13, 14, 31].

More recently, [12] proposes qualitative reasoning with description logic to describe consistency between geographic data sets. In [22] a set of abstract relations between entity classes is defined; and they could be used to discover redundancies and conflicts in sets of SICs. A proposal for fixing (changing) spatial database instances under different types of spatial inconsistencies is given in [29]. According to it, changes are applied over geometries in isolation; that is, they are not analyzed in combination with multiple SICs. In [27] some issues around query answering under violations of functional dependencies involving geometric attributes were raised. However, the problem of dealing with an inconsistent spatial database, while still obtaining meaningful answers, has not been systematically studied so far.

Consistent query answering (CQA) from inconsistent databases as a strategy of inconsistent tolerance has an extensive literature (cf. [4, 6, 10] for surveys). It was introduced and studied in the context of relational database in [2]. They defined consistent answers to queries as those that are invariant under all the minimal forms of restoring consistency of the original database. Thus, the notion of *repair* of an instance with respect to a set of ICs becomes a fundamental concept for defining consistent query answers. A *repair semantics* defines the admissible and consistent alternative instances to an inconsistent database at hand. More precisely, a *repair* of an inconsistent relational instance *D* is a consistent instance *D'* obtained from *D* by deleting or inserting whole tuples. The set of tuples by which *D* and *D'* differ is minimal under set

inclusion [2]. Other types of repair semantics have been studied in the relational case. For example, in [16, 32] repairs are obtained by allowing updates of attribute values in tuples.

In comparison to the relational case, spatial databases offer new alternatives and challenges when defining a repair semantics. This is due, in particular, to the use of complex attributes to represent geometries, their combination with thematic attributes, and the nature of spatial (topological) relations.

In this work we define a repair semantics for spatial databases with respect to a subset of spatial semantic integrity constraints (a.k.a. topo-semantic integrity constraints) [29], which impose semantic restrictions on topological predicates and combinations thereof. In particular, we treat spatial semantic integrity constraints that can be expressed by denials constraints. For example, they can specify that “two land parcels cannot internally intersect”. This class of constraints are neither standardized nor integrated into current spatial database management systems (DBMSs); they rather depend on the application, and must be defined and handled by the database developers. They are very important because they capture the semantics of the intended models. Spatial semantic integrity constraints will be simply called *spatial integrity constraints* (SICs). Other spatial integrity constraints [11] are *domain (topological or geometric) constraints*, and they refer to the geometry, topology, and spatial relations of the spatial data types. One of them could specify that “polygons must be closed”. Many of these geometric constraints are now commonly integrated into spatial DBMSs [23].

A definition of a repair semantics for spatial DBs and CQA for spatial range queries was first proposed in [28], where we discussed the idea of shrinking geometries to solve conflicting tuples and applied to CQA for range queries. In this paper we complement and extend our previous work with the following main contributions: (1) We formalize the repair semantics of a spatial database instance under violations of SICs. This is done through virtual changes of geometries that participate in violations of SICs. Unlike [28], we identify the admissible local transformations and we use them to provide an inductive definition of database repair. (2) Based on this formalization, a consistent answer to a spatial query is defined as an answer obtained from all the admissible repairs. Extending the results in [28], we now define CQA not only for range but also for spatial join queries. (3) Although the repair semantics and consistent query answers can be defined for a fairly broad class of SICs and queries, as it becomes clear soon, naive algorithms for computing consistent answers on the basis of the computation of all repairs are of exponential time. For this reason, CQA for a relevant subset of SICs and range and join queries is done via a *core computation*. This amounts to querying directly the intersection of all repairs of an inconsistent database instance, but without actually computing the repairs. We show cases where this core can be specified as a view of the original, inconsistent database. (4) We present an experimental evaluation with real and synthetic data sets that compares the cost of CQA with the cost of evaluating queries directly over the inconsistent database (i.e., ignoring inconsistencies).

The rest of the paper is organized as follows. In Section 2 we describe the spatial data model upon which we define the repair semantics and consistent query answers. A formal definition of repair for spatial inconsistent databases under SICs is introduced in Section 3. In Section 4 we define consistent answers to conjunctive queries. We analyze in particular the cases of range and join queries with respect to their computational properties. This leads us, in Section 5, to propose polynomial time algorithms (in data complexity) for consistent query answering with respect to a relevant class of SICs and queries. An experimental evaluation of the cost of CQA is provided in Section 6. Final conclusions and future research directions are given in Section 7.

2 Preliminaries

Current models of spatial database are typically seen as extensions of the relational data model (known as extended-relational or object-relational models) with the definition of abstract data types to specify spatial attributes. We now introduce a general spatio-relational database model that includes spatio-relational predicates (they could also be purely relational) and spatial ICs. It uses some of the definitions introduced in [25]. The model is independent of the geometric data model (e.g. Spaghetti [30], topological [18, 30], raster [19], or polynomial model [24]) underlying the representation of spatial data types.

A *spatio-relational database schema* is of the form $\Sigma = (\mathcal{U}, \mathcal{A}, \mathcal{R}, \mathcal{T}, \mathcal{O}, \mathcal{B})$, where: (a) \mathcal{U} is the possibly infinite database domain of atomic thematic values. (b) \mathcal{A} is a set of thematic, non-spatial, attributes. (c) \mathcal{R} is a finite set of spatio-relational predicates whose attributes belong to \mathcal{A} or are spatial attributes. Spatial attributes take admissible values in $\mathcal{P}(\mathbb{R}^m)$, the power set of \mathbb{R}^m , for an m that depends on the dimension of the spatial attribute. (d) \mathcal{T} is a fixed set of binary spatial predicates, with a built-in interpretation. (e) \mathcal{O} is a fixed set of geometric operators that take spatial arguments, also with a built-in interpretation. (f) \mathcal{B} is a fixed set of built-in relational predicates, like comparison predicates, e.g. $<, >, =, \neq$, which apply to thematic attribute values.

Each database predicate $R \in \mathcal{R}$ has a type $\tau(R) = [n, m]$, with $n, m \in \mathbb{N}$, indicating the number n of thematic attributes, and the spatial dimension m of the single spatial attribute (it takes values in $\mathcal{P}(\mathbb{R}^m)$).¹ In Example 1, $\tau(LandP) = [1, 2]$, since it has one thematic attribute (*idl*) and one spatial attribute (*geometry*) defined by a 2D polygon. In this work we assume that each relation R has a key of the form (1) formed by thematic attributes only:

$$\forall \bar{x}_1 \bar{x}_2 \bar{x}_3 s_1 s_2 (R(\bar{x}_1, \bar{x}_2; s_1) \wedge R(\bar{x}_1, \bar{x}_3; s_2) \rightarrow (\bar{x}_2 = \bar{x}_3 \wedge s_1 = s_2)), \quad (1)$$

where the \bar{x}_i are sequences of distinct variables representing thematic attributes of R , and the s_i are variables for geometric attributes. Here $s_1 = s_2$ means geometric equality; that is, the identity of two geometries.

A database instance D of a spatio-relational schema Σ is a finite collection of ground atoms (or *spatial database tuples*) of the form $R(c_1, \dots, c_n; s)$, where $R \in \mathcal{R}$, $\langle c_1, \dots, c_n \rangle \in \mathcal{U}^n$ contains the thematic attribute values, and $s \in Ad \subseteq \mathcal{P}(\mathbb{R}^m)$, where Ad is the class of admissible geometries (cf. below). The extension in a particular instance of a spatio-relational predicate is a subset of $\mathcal{U}^n \times Ad$. For simplicity, and to fix ideas, we will consider the case where $m = 2$.

Among the different abstraction mechanisms for modelling single spatial objects, we concentrate on regions for modelling real objects that have an extent. They are useful in a broad class of applications in Geographic Information Systems (GISs). More specifically, our model will be compatible with the specification of spatial operators (i.e., spatial relations or geometric operations) as found in current spatial DBMSs [23]. Following current implementations of DBMSs, regions could be defined as finite sets of polygons that, in their turn, are defined through their vertices. This would make regions finitely representable. However, in this work geometries will be treated at a more abstract level, which is independent of the spatial model used for geometric representation. In consequence, an admissible geometry of the Euclidean plane is either the empty geometry, g_\emptyset , which corresponds to the empty subset of the plane, or is a closed and bounded region with a positive area. It holds $g_\emptyset \cap g = g \cap g_\emptyset = g_\emptyset$, for every region g . From now on, empty geometries and regions of \mathbb{R}^2 are called *admissible geometries* and they form the class Ad .

Geometric attributes are complex data types, and their manipulation may have an important effect on the computational cost of certain algorithms and algorithmic problems. As usual, we are interested in *data*

¹For simplicity, we use one spatial attribute, but it is not difficult to consider a greater number of spatial attributes.

complexity, i.e., in terms of the size of the database. The *size* of a spatio-relational database can be defined as a function of the number of tuples and the representation size of geometries in those tuples.

We concentrate on binary (i.e., two-ary) spatial predicates that represent topological relations between *regions*. They have a fixed semantics, and become the elements of \mathcal{T} . There are eight basic binary relations over regions of \mathbb{R}^2 : *Overlaps* (*OV*), *Equals* (*EQ*), *CoveredBy* (*CB*), *Inside* (*IS*), *Covers* (*CV*), *Includes* (*IC*), *Touches* (*TO*), and *Disjoint* (*DJ*) [15, 26].² The semantics of the topological relations follows the point-set topology defined in [15], which is not defined for empty geometries. We will apply this semantics to our non-empty admissible geometries. For the case of the empty set, a separate definition will be given below. According to [15], an atom $T(x, y)$ becomes true if four conditions are simultaneously true. Those conditions are expressed in terms of emptiness (\emptyset) and non-emptiness ($\neg\emptyset$) of the intersection of their boundaries (δ) and interiors (\circ). The definitions can be found in Table ???. For example, for non-empty regions x, y , $TO(x, y)$ is true iff all of $\delta(x) \cap \delta(y) \neq \emptyset$, $\circ(x) \cap \circ(y) = \emptyset$, $\delta(x) \cap \circ(y) = \emptyset$, and $\circ(x) \cap \delta(y) = \emptyset$ simultaneously hold.

Relation	$\delta(x) \cap \delta(y)$	$\circ(x) \cap \circ(y)$	$\delta(x) \cap \circ(y)$	$\circ(x) \cap \delta(y)$
DJ(x,y)	\emptyset	\emptyset	\emptyset	\emptyset
TO(x,y)	$\neg\emptyset$	\emptyset	\emptyset	\emptyset
EQ(x,y)	$\neg\emptyset$	$\neg\emptyset$	\emptyset	\emptyset
IS(x,y)	\emptyset	$\neg\emptyset$	$\neg\emptyset$	\emptyset
CB(x,y)	$\neg\emptyset$	$\neg\emptyset$	$\neg\emptyset$	\emptyset
IC(x,y)	\emptyset	$\neg\emptyset$	\emptyset	$\neg\emptyset$
CV(x,y)	$\neg\emptyset$	$\neg\emptyset$	\emptyset	$\neg\emptyset$
OV(x,y)	$\neg\emptyset$	$\neg\emptyset$	$\neg\emptyset$	$\neg\emptyset$

Table 1: Definition of topological relations between regions based on point-set topology

In this work we exclude the topological relation *Disjoint* from \mathcal{T} . This decision is discussed in Section 3, where we introduce the repair semantics. In addition to the basic topological relations, we consider three *derived relations* that exist in current SQL languages and can be logically defined in terms of the other basic predicates: *Intersects* (*IT*), *Within* (*WI*), and *Contains* (*CO*). We also introduce a fourth relation, *Intersects (II)*, that holds when the interiors of two geometries intersect. It can be logically defined as the disjunction of *Overlaps*, *Within* and *Contains* (cf. Figure 2). For all the topological relations in \mathcal{T} , their converse (inverse) relation is within the set. Some of them are symmetric, like *Equals*, *Touches*, and *Overlaps*. For the non-symmetric relations, the converse relation of *CoveredBy* is *Covers*, of *Inside* is *Includes*, and of *Within* is *Contains*.

As mentioned before, the formal definitions of the topological relations [15, 26] do not consider the empty geometry as an argument. Indeed, at the best of our knowledge, no clear semantics for topological predicates with empty geometries exists. However, in our case we extend the definitions in order to deal with this case. This will allow us to use a classical bi-valued logic, where atoms are always true or false, but never undefined. According to our extended definition, for any $T \in \mathcal{T}$, $T(g_1, g_2)$ is false if $g_1 = g_\emptyset$ or $g_2 = g_\emptyset$. In particular, $IS(g, g_\emptyset)$ is false, for every admissible region g . In order to make comparisons with the empty region, we will introduce and use a special predicate $IsEmpty(\cdot)$ on admissible geometries, such

²The names of relations chosen here are in agreement with the names used in current SQL languages [23], but differ slightly from the names found in the research literature. The relations found in current SQL languages are represented in Figure 2 with thick borders.




Figure 2: Subsumption lattice of topological relations between regions: OV (Overlaps), CB (CoveredBy), IS (Inside), EQ (Equals), CV (Covers), IC (Includes), TO (Touches), DJ (Disjoint), IT (Intersects), II (Intersects), WI (Within), and CO (Contains).

that $IsEmpty(s)$ is true iff $s = g_{\emptyset}$.

Notice that the semantics of the topological predicates, even for non-empty regions, may differ from the intuitive set-theoretic semantics one could assign to them. For example, for an admissible and non-empty geometry g , $OV(g, g)$ is false (due to the conditions in the last two columns in Table ??). In consequence, the constraint $\forall x \forall s \neq g_{\emptyset} \neg (R(x; s) \wedge OV(s, s))$ is satisfied.

Given a database instance, additional spatial information is usually computed from the explicit geometric data by means of the spatial operators in \mathcal{O} associated with Σ . Some relevant operators are: *Area*, *Union* (binary), *Intersection*, *Difference*, *Buffer*, and *Union Aggregation* (*GeomUnion*).³ (Cf. [23] for the complete set of spatial predicates defined within the Open GIS Consortium.) There are several spatial operators used in this work; however, we will identify a particular subset \mathcal{O}^a of spatial operators in \mathcal{O} , i.e., $\mathcal{O}^a \subseteq \mathcal{O}$, which will be defined for all admissible geometries and used to shrink geometries with the purpose of restoring consistency, as we describe in Section 3.

Definition 1 The set \mathcal{O}^a of *admissible operations* contains the following geometric operations on admissible geometries g and g' :

- (1) *Difference*(g, g') is the topological closure of the set-difference.
- (2) *Buffer*(g, d) is the geometry obtained by buffering a distance d around g , where d is a distance unit. *Buffer*(g, d) returns a closed region \bar{g} containing geometry g , such that every point in the boundary of \bar{g} is at a distance d from some point of the boundary of g . In particular, $Buffer(g_{\emptyset}, d) = g_{\emptyset}$. \square

Notice that these operators, when applied to admissible geometries, produce admissible geometries.

Remark 1 The value of d in Definition 1 is instance dependent. It should be precomputed from the spatial input data. For this work, we consider d to be a fixed value associated with the minimum distance between geometries in the cartographic scale of the database instance. \square

³Operator *GeomUnion* returns the geometry that represents the point set union of all geometries in a given set, an operator also known as a spatial aggregation operator. Although this function is part of SQL for several spatial databases (Postgres/PostGIS, Oracle), it is not explicitly defined in the OGC specification [23].




Figure 3: A spatial database instance

A schema Σ determines a many-sorted, first-order (FO) language $\mathcal{L}(\Sigma)$ of predicate logic. It can be used to syntactically characterize and express SICs. For simplicity, we concentrate on *denial SICs*,⁴ which are sentences of the form:

$$\forall^{\neq g_0} \bar{s} \forall \bar{x} \neg \left(\bigwedge_{i=1}^m R_i(\bar{x}_i; s_i) \wedge \varphi \wedge \bigwedge_{j=1}^n T_j(v_j, w_j) \right). \quad (2)$$

Here, $\bar{s} = s_1 \cdots s_m$, $\bar{x} = \bar{x}_1 \cdots \bar{x}_m$ are finite sequences of geometric and thematic variables, respectively, and $0 < m, n \in \mathbb{N}$. Thus, each \bar{x}_i is a finite tuple of thematic variables and will be treated as a set of attributes, such that $\bar{x}_i \subseteq \bar{x}_j$ means that the variables in \bar{x}_i are also variables in \bar{x}_j . Also, $\forall \bar{x}$ stands for $\forall x_1 \cdots \forall x_m$; and $\forall^{\neq g_0} \bar{s}$ stands for $\forall s_1 \cdots \forall s_m$, with the universal quantifiers ranging over all the non-empty admissible geometries (i.e. regions). Here, $v_j, w_j \in \bar{s}$, $R_1, \dots, R_m \in \mathcal{R}$, φ is a formula containing built-in atoms over thematic attributes, and $T_j \in \mathcal{T}$. A constraint of the form (2) prohibits certain combinations of database atoms. Since topological predicates for empty geometries are always false, the restricted quantification over non-empty geometries in the constraints could be eliminated. However, we do not want to make the satisfaction of the constraints rely on our particular definition of the topological predicates for the empty region. In this way, our framework becomes more general, robust and modular, in the sense that it would be possible to redefine the topological predicates for the empty region without affecting our approach and results.

Example 2 Figure 3 shows an instance for the schema $\mathcal{R} = \{LandP(idl, name, owner; geometry), Building(idb; geometry)\}$. Dark rectangles represent buildings and white rectangles represent land parcels. In *LandP*, the thematic attributes are *idl*, *name* and *owner*, whereas *geometry* is the spatial attribute of dimension 2. Similarly for *Building*, which has only *idb* as a thematic attribute.

The following sentences are denial SICs: (The symbol $\bar{\forall}$ stands for the universal closure of the formula that follows it.)

$$\bar{\forall} \neg (LandP(idl_1, n_1, o_1; s_1) \wedge LandP(idl_2, n_2, o_2; s_2) \wedge idl_1 \neq idl_2 \wedge IIntersects(s_1, s_2)). \quad (3)$$

$$\bar{\forall} \neg (Building(idb; s_1) \wedge LandP(idl, n, o; s_2) \wedge Overlaps(s_1, s_2)). \quad (4)$$

The SIC (3) says that geometries of land parcels with different *ids* cannot internally intersect (i.e., they can only be disjoint or touch). The SIC (4) establishes that building blocks cannot (partially) overlap land parcels. \square

⁴Denial constraints are easier to handle in the relational case as consistency with respect to them is achieved by tuple deletions only [6].

A database instance D for schema Σ can be seen as an interpretation structure for the language $\mathcal{L}(\Sigma)$. For a set Ψ of SICs in $\mathcal{L}(\Sigma)$, $D \models \Psi$ denotes that each of the constraints in Ψ is true in (or satisfied by) D . In this case, we say that D is *consistent* with respect to Ψ . Correspondingly, D is *inconsistent* with respect to Ψ , denoted $D \not\models \Psi$, when there is a $\psi \in \Psi$ that is *violated* by D , i.e., not satisfied by D . The instance in Example 2 is consistent with respect to its SICs.

In what follows, we will assume that the set Ψ of SICs under consideration is logically consistent; i.e., that there exists a non-empty database instance D (not necessarily the one at hand), such that $D \models \Psi$. For example, any set of SICs containing a constraint of the form $\forall^{g \neq g_0} s \forall \bar{x} \neg (R(\bar{x}; s) \wedge Equals(s, s))$ is logically inconsistent. The analysis of whether a set of SICs is logically consistent or not is out of the scope of this work.

3 A Repair Semantics

Different alternatives for update-based consistency restoration of spatial databases are discussed in [28]. One of the key criteria to decide about the update to apply is minimality of geometric changes. Another important criteria may be the semantics of spatial objects, which makes changes over the geometry of one type of object more appropriate than others. For this work, the repair semantics is a rule applied automatically. It assumes that no previous knowledge about the quality and relevance of geometries exists and, therefore, it assumes that geometries are all equally important.

On the basis on the minimality condition on geometric changes and the monotonicity property of some topological predicates [28], we propose to solve inconsistencies with respect to SICs of the form (2) through shrinking of geometries. Notice that this repair semantics will be used as an instrumental concept to formalize consistent query answers (no actual modification over the database occurs). As such, it defines what part of the geometry is not in conflict with respect to a set of integrity constraints and can, therefore, be part of a consistent answer.

Shrinking geometries eliminates conflicting parts of geometries without adding new uncertain geometries by enlargement. In this way, we are considering a proper subset of the possible changes to fix spatial databases proposed in [29]. We disregard translating objects, because they will carry potentially new conflicts; and also creating new objects (object splitting), because we would have to deal with null or unknown thematic attributes.

The SICs of the form (2) exclude the topological predicate *Disjoint*. The reason is that falsifying an atom $DJ(g_1, g_2)$ by shrinking geometries is not possible, unless we make one of them empty. However, doing so would heavily depend upon our definition of this topological predicate for empty regions. Since we opted for not making our approach and results depend on this particular definition, we prefer to exclude the *Disjoint* predicate from our considerations. The study of other repair semantics that sensibly includes the topological predicate *Disjoint* will be left for future work.

Technically, a database D violates a constraint $\forall \bar{x}_1 \bar{x}_2 \forall^{g \neq g_0} s_1 s_2 \neg (R_1(\bar{x}_1; s_1) \wedge R_2(\bar{x}_2; s_2) \wedge \varphi \wedge T(s_1, s_2))$, with $T \in \mathcal{T}$,⁵ when there are data values $\bar{a}_1, \bar{a}_2, g_1, g_2$, with $g_1, g_2 \neq g_0$, for the variables in the constraint such that $(R_1(\bar{x}_1; s_1) \wedge R_2(\bar{x}_2; s_2) \wedge \varphi \wedge T(s_1, s_2))$ becomes true in the database under those values. This is denoted with $D \models (R_1(\bar{x}_1; s_1) \wedge R_2(\bar{x}_2; s_2) \wedge \varphi \wedge T(s_1, s_2)) [\bar{a}_1, \bar{a}_2, g_1, g_2]$. When this is the case, it is possible to restore consistency of D by shrinking g_1 or g_2 such that $T(g_1, g_2)$ becomes false.

⁵For simplicity and without lost of generality, in the examples we consider denial constraints with at most two spatio-relational predicates and one topological predicate. However, a denial constraint of the form (2) may have more spatio-relational predicates and topological predicates.

We can compare geometries, usually an original geometry and its shrunk version, by means of a distance function that refers to their areas. We assume that $area \in \mathcal{O}$ is an operator that computes the area of a geometry.

Definition 2 For regions g_1, g_2 , $\delta(g_1, g_2) = area(Difference(g_1, g_2) \cup Difference(g_2, g_1))$. \square

Since we will compare a region g_1 with a region g_2 obtained by shrinking g_1 , it will hold $\delta(g_1, g_2) \geq 0$. Indeed, when comparing $g_2 \subseteq g_1$ ⁶, the distance function can be simplified by $\delta(g_1, g_2) = area(Difference(g_1, g_2))$. We will assume that it is possible to compare geometries through the distance function by correlating their tuples, one by one. This requires a correspondence between instances.

Definition 3 Let D, D' be database instances of schema Σ . D' is $(D, f_{D'})$ -indexed if $f_{D'}$ is a bijective function from D to D' , such that, for all c_1, \dots, c_n, s : $f_{D'}(R(c_1, \dots, c_n; s)) = R(c_1, \dots, c_n; s')$, for some region s' . \square

In a $(D, f_{D'})$ -indexed instance D' we can compare tuples one by one with their counterparts in instance D . In particular, we can see how the geometric attribute values differ. In some cases there is an obvious function $f_{D'}$, for example, when there is a key from a subset of \mathcal{A} to the spatial attribute S , or when relations have a surrogate key for identification of tuples. In these cases we simply use the notion of D -indexed. When the context is clear, we also use f instead of $f_{D'}$.

Example 3 (example 2 cont.) Consider the relational schema $LandP(idl, name, owner; geometry)$. For the instance D given in Example 2, the following instance D' is (D, f) -indexed

LandP			
<i>idl</i>	<i>name</i>	<i>owner</i>	<i>geometry</i>
<i>idl</i> ₁	<i>n</i> ₁	<i>o</i> ₁	<i>g</i> ₇
<i>idl</i> ₂	<i>n</i> ₂	<i>o</i> ₂	<i>g</i> ₈
<i>idl</i> ₃	<i>n</i> ₃	<i>o</i> ₃	<i>g</i> ₉

Here, $f(LandP(idl_1, n_1, a_1; g_1)) = LandP(idl_1, n_1, a_1; g_7)$, etc. \square

When restoring consistency, it may be necessary to consider different combinations of tuples and SICs. Eventually, we should obtain a new instance, hopefully consistent, that we have to compare to the original instance in terms of their distance.

Definition 4 Let D, D' be spatial database instances over the same schema Σ , with D' (D, f) -indexed. The distance $\Delta(D, D')$ between D and D' is the numerical value $\Delta(D, D') = \sum_{\bar{t} \in D} \delta(\Pi_S(\bar{t}), \Pi_S(f(\bar{t})))$, where $\Pi_S(\bar{t})$ is the projection of tuple \bar{t} on its spatial attribute S . \square

Now it is possible to define a “repair semantics”, which is independent of the geometric operators used to shrink geometries.

Definition 5 Let D be a spatial database instance over schema Σ , Ψ a set of SICs, such that $D \not\models \Psi$. (a) An s -repair of D with respect to Ψ is a database instance D' over Σ , such that: (i) $D' \models \Psi$. (ii) D' is (D, f) -indexed. (iii) For every tuple $R(c_1, \dots, c_n; g) \in D$, if $f(R(c_1, \dots, c_n; g)) = R(c_1, \dots, c_n; g')$, then $g' \subseteq g$. (b) A minimal s -repair D' of D is a repair of D such that, for every repair D'' of D , it holds $\Delta(D, D'') \geq \Delta(D, D')$. \square

⁶ \subseteq stands for geometric inclusion

Proposition 1 If D is consistent with respect to Ψ , then D is also its only minimal s-repair.

Proof: For $D' = D$, it holds: (i) $D' \models \Psi$, (ii) D' is (D, f) -indexed, (iii) for every tuple $R(c_1, \dots, c_n; g) \in D$, if $f(R(c_1, \dots, c_n; g)) = R(c_1, \dots, c_n; g')$, then $g' = g$. In this case, $\Delta(D, D') = 0$. Any other consistent instance D'' obtained by shrinking any of D 's geometries and still obtaining admissible geometries gives $\Delta(D, D'') > 0$. \square

This is an “ideal and natural” repair semantics that defines a collection of *semantic repairs*. The definition is purely set-theoretic and topological in essence. It is worth exploring the properties of this semantics and its impact on properties of consistent query answers (as invariant under minimal s-repairs) and on logical reasoning about them. However, for a given database instance we may have a continuum and infinite number of s-repairs since between two points we have an infinite number of points, which we want to avoid for representational and computational reasons.

In this work we will consider an alternative repair semantics that is more operational in nature (cf. Definition 8), leaving the previous one for reference. This *operational* definition of repair makes it possible to deal with repairs in current spatial DBMSs and in terms of standard geometric operators (cf. Lemma 1). Under this definition, there will always be a finite number of repairs for a given instance. Consistency will be restored by applying a finite sequence of admissible transformation operations to conflicting geometries.

It is easy to see that each true relationship (atom) of the form $T(g_1, g_2)$, with $T \in \mathcal{T}$, can be falsified by applying an admissible transformation in \mathcal{O}^a to g_1 or g_2 . Actually, they can be falsified in a *canonical way*. These canonical falsification operations for the different topological atoms are presented in Table 1. They have the advantages of: (a) being defined in terms of the admissible operators, (b) capturing the repair process in terms of the elimination of conflicting parts of geometries, and (c) changing one of the geometries participating in a conflict.

More specifically, in Table 1 we indicate, for each relation $T \in \mathcal{T}$, alternative operations that falsify a true atom of the form $T(g_1, g_2)$. Each of them makes changes on one of the geometries, leaving the other geometry unchanged. The list of *canonical transformations* in this table prescribes particular ways of applying the admissible operators of Definition 1. Later on, they will also become the *admissible or legal* ways of transforming geometries with the purpose of restoring consistency.

For example, Table 1 shows that for *Overlaps*(*OV*), there are in principle four ways to make false an atom *Overlaps*(g_1, g_2) that is true. These are the alternatives 1. to 4. in that entry, where alternatives 1. and 2. change geometry g_1 ; and alternatives 3. and 4. change geometry g_2 . Only one of these alternatives that satisfies its condition is expected to be chosen to falsify the atom. A minimal way to change a geometry depends on the relative size between overlapping and non-overlapping areas: (i) when the overlapping area between g_1 and g_2 is smaller than or equal to their non-overlapping areas, a minimal change over geometry g_1 is *Difference*(g_1, g_2), and over g_2 is *Difference*(g_2, g_1) (cases 1. and 3. for *OV* in Table 1). (ii) When the non-overlapping areas of g_1 or g_2 are smaller than the overlapping area, a minimal change over geometry g_1 is *Difference*($g_1, \text{Difference}(g_1, g_2)$), and over geometry g_2 is *Difference*($g_2, \text{Difference}(g_2, g_1)$) (cases 2. and 4. for *OV* in Table 1).

For the case when *Equals*(g_1, g_2) is true, the transformations in Table 1 make either geometry, g_1 or g_2 empty to falsify the atom. However, there are other alternatives that by shrinking geometries would achieve the same result, but also producing smaller changes in terms of the affected area. A natural candidate update consists in applying the transformation $g'_1 = \text{Difference}(g_1, \text{Buffer}(\text{Boundary}(g_2), d))$ (similarly and alternatively for g_2). In this case, we just take away from g_1 the part of the internal area of width d surrounding the boundary of g_1 , to make it different from g_2 . We did not follow this alternative for practical reasons: having two geometries that are topologically equal could, in many cases, be the result of duplicate data, and

Pred. T	A true atom $T(g_1, g_2)$ becomes a false atom $T(g'_1, g'_2)$ with
<i>OV</i>	<ol style="list-style-type: none"> 1. If $area(g_1 \cap g_2) \leq area(g_1 \setminus g_2)$: $g'_1 = Difference(g_1, g_2)$, $g'_2 = g_2$. 2. If $area(g_1 \cap g_2) > area(g_1 \setminus g_2)$: $g'_1 = Difference(g_1, Difference(g_1, g_2))$, $g'_2 = g_2$. 3. If $area(g_1 \cap g_2) \leq area(g_2 \setminus g_1)$: $g'_2 = Difference(g_2, g_1)$, $g'_1 = g_1$. 4. If $area(g_1 \cap g_2) > area(g_2 \setminus g_1)$: $g'_2 = Difference(g_2, Difference(g_2, g_1))$, $g'_1 = g_1$.
<i>IS, CB</i>	<ol style="list-style-type: none"> 1. If $area(g_1 \cap g_2) \leq area(g_2 \setminus g_1)$: $g'_2 = Difference(g_2, g_1)$, $g'_1 = g_1$. 2. If $area(g_1 \cap g_2) > area(g_2 \setminus g_1)$: $g'_2 = Difference(g_2, Difference(g_2, g_1))$, $g'_1 = g_1$. 3. $g'_1 = Difference(g_1, g_2)$, $g'_2 = g_2$.
<i>IC, CV</i>	<ol style="list-style-type: none"> 1. If $area(g_1 \cap g_2) \leq area(g_1 \setminus g_2)$: $g'_1 = Difference(g_1, g_2)$, $g'_2 = g_2$. 2. If $area(g_1 \cap g_2) > area(g_1 \setminus g_2)$: $g'_1 = Difference(g_1, Difference(g_1, g_2))$, $g'_2 = g_2$. 3. $g'_2 = Difference(g_2, g_1)$, $g'_1 = g_1$.
<i>II, WI, CO</i>	<ol style="list-style-type: none"> 1. $g'_1 = Difference(g_1, g_2)$, $g'_2 = g_2$. 2. $g'_2 = Difference(g_2, g_1)$, $g'_1 = g_1$.
<i>TO, IT</i>	<ol style="list-style-type: none"> 1. $g'_1 = Difference(g_1, buffer(g_2, d))$, $g'_2 = g_2$. 2. $g'_2 = Difference(g_2, buffer(g_1, d))$, $g'_1 = g_1$. <p>(See Remark ?? for definition of d)</p>
<i>EQ</i>	<ol style="list-style-type: none"> 1. $g'_1 = g_{\emptyset}$, $g'_2 = g_2$. 2. $g'_2 = g_{\emptyset}$, $g'_1 = g_1$.

Table 2: Admissible transformations

one of them should be eliminated. Moreover, this alternative, in comparison with the officially adopted in this work, may create new conflicts with respect to other SICs. Avoiding them whenever possible will be used later, when designing a polynomial algorithm for CQA based on the core of an inconsistent database instance (see Section 5).

Table 1 shows that *Touches* and *Intersects* are predicates for which the eliminated area is not completely delimited by the real boundary of objects. Actually, we need to separate the touching boundaries. We do so

by buffering a distance d around one of the geometries and taking the overlapping part from the other one.⁷

The following result is obtained directly from Table 1.

Lemma 1 For each topological predicate $T \in \mathcal{T}$ and true ground atom $T(g_1, g_2)$, there are geometries g'_1, g'_2 obtained by means of the corresponding admissible transformation in Table 1, such that $T(g'_1, g'_2)$ becomes false. \square

The following definition defines, for each geometric predicate T , a binary geometric operator tr^T such that, if $T(g_1, g_2)$ is true, then $tr^T(g_1, g_2)$ returns a geometry g'_1 such that $T(g'_1, g_2)$ becomes false. The definition is based on the transformations that affect geometry g_1 in Table 1.

Definition 6 Let $T \in \mathcal{T}$ be a topological predicate. We define an admissible transformation operator $tr^T : Ad \times Ad \rightarrow Ad$ as follows:

(a) If $T(g_1, g_2)$ is false, then $tr^T(g_1, g_2) := g_1$.

(b) If $T(g_1, g_2)$ is true, then:

$$tr^T(g_1, g_2) := \begin{cases} \text{Difference}(g_1, g_2) & \text{if } \text{area}(g_1 \cap g_2) \leq \text{area}(g_1 \setminus g_2) \\ \text{Difference}(g_1, \text{Difference}(g_1, g_2)) & \text{otherwise} \end{cases}$$

for $T \in \{OV, IC, CV\}$;

$$tr^T(g_1, g_2) := \text{Difference}(g_1, g_2), \text{ for } T \in \{IS, CB, II, WI, CO\};$$

$$tr^T(g_1, g_2) := \text{Difference}(g_1, \text{Buffer}(g_2, d)) \text{ for } T \in \{TO, IT\};$$

$$tr^T(g_1, g_2) := g_\emptyset \text{ for } T \in \{EQ\}. \quad \square$$

It can be easily verified that the admissible operations tr^T , applied to admissible geometries, produce admissible geometries. They can be seen as macros defined in terms of the basic operations in Definition 1, and inspired by Table 1. The idea is that the operator tr^T takes (g_1, g_2) , for which $T(g_1, g_2)$ is true, and makes the latter false by transforming g_1 into g'_1 , i.e., $T(g'_1, g_2)$ becomes false.

Definition ?? can also be used to formalize the transformations on geometry g_2 indicated in Table 1. First, notice that for the converse predicate T^c of predicate T it holds: $T^c(g_1, g_2)$ true iff $T(g_2, g_1)$. Secondly, the converse of a transformation operator can be defined by $(tr^T)^c := tr^{(T^c)}$. In consequence, we can apply tr^{T^c} to (g_2, g_1) , obtaining the desired transformation of geometry g_2 . In this way, all the cases in Table 1 are covered. For example, if we want to make false a true atom $Inside(g_1, g_2)$, we can apply $tr^{IS}(g_1, g_2)$, but also $tr^{IC}(g_2, g_1)$.

Example 4 Table 2 illustrates the application of the admissible transformations to restore consistency of predicates $T \in \{Overlaps, Touches\}$. The dashed boundary is the result of applying $Buffer(g, d)$. \square

We now define the notion of *accessible instance* that results from an original instance, after applying admissible transformation operations to geometries. The application of sequences of operators solves violations of SICs. Accordingly, the accessible instances are defined by induction.

⁷The buffer operator does not introduce new points in the geometric representation of objects, but it translates the boundary a distance d outwards.

$T(g_1, g_2)$	Original	$Tr^T(g_1, g_2)$	$Tr^T(g_2, g_1)$
OV			
TO			

Table 3: Examples of admissible transformations

Definition 7 Let D be a database instance. D' is an *accessible instance* from D (with respect to a finite set of SICs Ψ), if D' is obtained after applying, a finite number of times, the following inductive rules (any of them, when applicable):

- (1). $D' = D$.
- (2). There is an accessible instance D_0 from D , such that, for some $\psi \in \Psi$ with a topological predicate T , $D_0 \not\models \psi^8$ through tuples $R_1(\bar{a}_1; g_1)$ and $R_2(\bar{a}_2; g_2)$ in D_0 , for which $T(g_1, g_2)$ is true; and
 - (a) $D' = D_0 \setminus \{R_1(\bar{a}_1, g_1)\} \cup \{R_1(\bar{a}_1, tr^T(g_1, g_2))\}$, or
 - (b) $D' = D_0 \setminus \{R_2(\bar{a}_2, g_2)\} \cup \{R_2(\bar{a}_2, tr^{T^c}(g_2, g_1))\}$. □

Example 5 Consider the database instance in Figure 4(a) that is inconsistent with respect to SIC (3). An accessible instance from this inconsistent database is in Figure 4(b), where only g_1 has changed. This can be expressed in the following way: $LandP(idl_1, n_1, o_1; g'_1) = LandP(idl_1, n_1, o_1; tr^{II}(tr^{II}(g_1, g_2), g_3))$. □

Given a database D , possibly inconsistent, we are interested in those accessible instances D' that are consistent, i.e., $D' \models \Psi$. Even more, having the repairs in mind, we have to make sure that admissible instances from D can still be indexed with D .


Proposition 2 Let D' be an accessible instance from D . Then, D' is f -indexed to D via an index function f , that can be defined by induction on D' .

Proof: To simplify the presentation, we will assume that D has an index (or surrogate key) i_0 , that is a one-to-one mapping from D to an initial segment $[1, N]$ of \mathbb{N} . Let D' be an accessible instance from D . We define $i_{D'}(R(\bar{a}; g)) \in \mathbb{N}$ for tuples in D' by induction on D' :

- (1). If $D' = D$ and $R(\bar{a}; g) \in D$, $i_{D'}(R(\bar{a}; g)) = i_0(R(\bar{a}; g))$.
- (2). If there is an accessible instance D_0 from D and $D_0 \not\models \psi \in \Psi$ through the atoms $R_1(\bar{a}_1; g_1)$, $R_2(\bar{a}_2; g_2)$, and $T(g_1, g_2)$ with T and T^c the converse relation of T :


⁸ ψ may have more than one topological predicate.

LandP			
<i>idl</i>	<i>name</i>	<i>owner</i>	<i>geometry</i>
idl_1	n_1	o_1	g_1
idl_2	n_2	o_2	g_2
idl_3	n_3	o_3	g_3



(a)

LandP			
<i>idl</i>	<i>name</i>	<i>owner</i>	<i>geometry</i>
idl_1	n_1	o_1	g'_1
idl_2	n_2	o_2	g_2
idl_3	n_3	o_3	g_3



(b)

Figure 4: An accessible instance: (a) original instance, and (b) accessible transformation (geometry with thick boundary changed)

- (a) $D' = D_0 \setminus \{R_1(\bar{a}_1, g_1)\} \cup \{R_1(\bar{a}_1, tr^T(g_1, g_2))\}$, and $i_{D'}(R_1(\bar{a}_1; tr^T(g_1, g_2))) = i_{D_0}(R_1(\bar{a}_1, g_1))$ and $i_{D'}(R_2(\bar{a}_2; g_2)) = i_{D_0}(R_2(\bar{a}_2, g_2))$, or
(b) $D' = D_0 \setminus \{R_2(\bar{a}_2, g_2)\} \cup \{R_2(\bar{a}_2, tr^{Tc}(g_2, g_1))\}$, and $i_{D'}(R_1(\bar{a}_1; g_1)) = i_{D_0}(R_1(\bar{a}_1, g_1))$ and $i_{D'}(R_2(\bar{a}_2; tr^{Tc}(g_2, g_1))) = i_{D_0}(R_2(\bar{a}_2, g_2))$. \square

Any two accessible instances D' and D'' can be indexed via D in a natural way, and thus, they can be compared tuple by tuple. In the following, we will assume, when comparing any two accessible instances in this way, that there is such an underlying index function f . Now we give the definition of operational repair.

Definition 8 Let D be an instance over schema Σ and Ψ a finite set of SICs. (a) An *o-repair* of D with respect to Ψ is an instance D' that is accessible from D , such that $D' \models \Psi$. (b) A *minimal o-repair* D' of D is an o-repair of D such that, for every o-repair D'' of D , $\Delta(D, D'') \geq \Delta(D, D')$. (c) $Rep(D, \Psi)$ denotes the set of minimal o-repairs of D with respect to Ψ . \square

The distances $\Delta(D, D'')$ and $\Delta(D, D')$ in this definition are relative to the corresponding index functions, whose existence is guaranteed by Proposition 2. Unless otherwise stated, this is the repair semantics we refer to in the remainder of the paper, in particular, in the definition of consistent query answer in Section 4. In consequence, in the following a repair is an o-repair, and the same applies to minimal repairs. Even more, whenever we refer to repairs, we should understand that minimal repairs are intended.

Example 6 Consider database schema in Example 2. The instance D in Figure 5 is inconsistent with respect to the SICs (3) and (4), because the land parcels with geometries g_2 and g_3 overlap, and so do the land parcels with geometries g_2 and g_4 . Likewise, buildings with geometry g_5 and g_6 *partially* overlap land parcels with geometries g_1 and g_2 , respectively.

LandP			
<i>idl</i>	<i>name</i>	<i>owner</i>	<i>geometry</i>
<i>idl</i> ₁	<i>n</i> ₁	<i>o</i> ₁	<i>g</i> ₁
<i>idl</i> ₂	<i>n</i> ₂	<i>o</i> ₂	<i>g</i> ₂
<i>idl</i> ₃	<i>n</i> ₃	<i>o</i> ₃	<i>g</i> ₃
<i>idl</i> ₄	<i>n</i> ₄	<i>o</i> ₄	<i>g</i> ₄

Building	
<i>idb</i>	<i>geometry</i>
<i>idb</i> ₁	<i>g</i> ₅
<i>idb</i> ₂	<i>g</i> ₆




Figure 5: Inconsistent database instance




Figure 6: Minimal repairs

Figure 6 shows the two minimal repairs of D . In them, the regions with thicker boundaries are the regions that have their geometries changed. For the minimal repair in Figure 6(a), the inconsistency involving geometries g_2 and g_3 is repaired by applying $Difference(g_2, g_3)$ to g_2 , i.e., removing from g_2 the whole overlapping geometry, and keeping the geometry of g_3 as originally. Notice that due to the interaction between integrity constraints, if we apply $Difference(g_3, g_2)$ to g_3 , i.e., we remove the whole overlapping area from g_3 , we still have an inconsistency, because the building with geometry g_6 will continue partially overlapping geometries g_2 and g_3 . Thus, this change will require an additional transformation to ensure that g_6 is completely covered or inside of g_3 .

In the same minimal repair (Figure 6(a)), the inconsistency between g_2 and g_4 is repaired by shrinking g_2 , eliminating its area that overlaps g_4 . This is obtained by applying $Difference(g_2, g_4)$ to g_2 . Finally, the inconsistency between g_1 and g_5 is repaired by removing from g_5 its part that does not overlap with geometry g_1 . In principle, we could have repaired this inconsistency by eliminating the overlapping region between g_1 and g_5 , but this is not a minimal change.

In the second minimal repair (Figure 6(b)), geometries g_2 and g_5 undergo the same changes than those in the first minimal repair (Figure 6(a)), but the inconsistency between g_2 and g_4 is restored by eliminating geometry g_4 , i.e., applying $Difference(g_4, g_2) = g_\emptyset$. \square

Notice that, by applying admissible transformation operators to restore consistency, the whole part of a geometry that is in conflict with respect to another geometry is removed. In consequence, given that there are finitely many geometries in the database instance and finitely many SICs, a finite number of applications of admissible transformations are sufficient to restore consistency. This contrasts with the s-repair semantics, which can yield even a continuum of possible consistency-restoration transformations. Keeping the number of repairs finite may be crucial for certain mechanisms for computing consistent query answers, as those as we will show in the next sections. Actually, we will use existing geometric operators as implemented in spatial DBMSs in order to capture and compute the consistency-restoring geometric transformations. This will be eventually used to obtain consistent query answers for an interesting class of spatial queries and SICs in Section 5.2.

LandP			
<i>idl</i>	<i>name</i>	<i>owner</i>	<i>geometry</i>
idl_1	n_1	o_1	g_1
idl_2	n_2	o_2	g_2

Building	
<i>idb</i>	<i>geometry</i>
idb_1	g_3
idb_2	g_4




Figure 7: An inconsistent database instance




Figure 8: Minimal repairs: (a) minimal s-repair and (b) minimal o-repairs (thick boundaries show geometries that have changed)

Despite the advantages of using o-repairs, the following example shows that an o-repair may not be minimal under the s-repair semantics.

Example 7 The instance D in Figure 7 is inconsistent with respect to the SICs (3) and (4), because the land parcels with geometries g_1 and g_2 internally intersect and buildings with geometry g_3 and g_4 overlap land parcels with geometries g_1 and g_2 , respectively.

Figures 8(a) and (b) show the minimal s-repair (Definition 5) and o-repairs of D (Definition 8), respectively. In them, the regions with thicker boundaries are the regions that had their geometries changed. Here, by applying s-repair semantics we obtain one minimal repair (Figure 8(a)) that takes the partial conflicting parts from both land parcels g_1 and g_2 in conflict, and leave unchanged the geometries of buildings g_3 and g_4 . Instead, for the o-repair semantics, each repair takes the whole conflicting parts from one of the land parcels g_1 or g_2 in order to satisfy SIC (3), and to satisfy SIC (4), each repair eliminates the conflict between the new version of g_1 and building g_3 or between the new version of g_2 and building g_4 . This makes up to four possible o-repairs (Figure 8(b)), which are not minimal with respect to the single s-repair. \square

S-repairs may take away only parts of a geometry that participate in a conflict. On the other side, they do not force a conflicting geometry to become empty in cases where o-repairs would do so. For instance, consider a true atom $Equals(g_1, g_2)$ that has to be falsified. A s-repair can be obtained by shrinking one of the two geometries just a little, without making it empty. However, by using admissible transformations, we can only falsify this atom by making one of the geometries empty. In this case, a minimal o-repair is not a minimal s-repair.

Proposition 3 Let D be a database instance, and Ψ a set of SICs. Then the following properties for o-repairs hold: (a) If D is consistent with respect to Ψ , then D is its only minimal o-repair. (b) If D' is an (D, f) -indexed o-repair of D and $f(R(\bar{a}; g)) = R(\bar{a}; g')$, then $g' \subseteq g$. (c) The set of o-repair for D is finite and non-empty.

Proof: (a) By the inductive definition of o-repair, an admissible transformation operator is applied to a geometry g when g is in conflict with other geometry g' in D . Since a consistent database instance does not




Figure 9: Example of a region's boundary after geometric transformation

contain conflicting tuples, none of the transformations operators is applicable and the consistent database instance is its only o-repair.

(b) The application of each admissible transformation $tr^T(g_1, g_2)$, with $T \in \mathcal{T}$, has five possible outcomes: $g_1, g_\emptyset, Difference(g_1, g_2), Difference(g_1, Difference(g_1, g_2))$, or $Difference(g_1, Buffer(g_2, d))$. Then, by definition of operator *Difference* (cf. Definition 1), $tr^T(g_1, g_2) \subseteq g_1$.

(c) D has a finite number N of tuples; and Ψ , a finite number of integrity constraints. In consequence, there is a finite number of conflicts, i.e., sets of tuples that simultaneously participate in the violation of one element ψ of Ψ via their geometries. Each of these conflicts are solved by shrinking some of those geometries. Each application of an operator O , chosen for a finite set of them, according to the inductive definition of o-repair solves an existing conflict by falsifying at least one of the \mathcal{T} -atoms in a ground instance of ψ . In principle, the application of such an operator O may produce new conflicts; however, it strictly decreases the total geometrical area of the database instance. More precisely, if $A(D') := \sum_{R(\bar{t}, g) \in D'} area(g)$, then $A(D') > A(O(D'))$, where $O(D')$ is the instance resulting from the conflict resolving operator O to D' . In particular, A_0 is the area $A(D)$ of the original instance D .

Now we reason by induction on the structure of o-repairs. The application of a one-conflict solving operator O to an instance D_{n-1} produces an instance D_n with $A(D_n) < A(D_{n-1})$. Moreover, $A(D_{n-1}) - A(D_n) > \epsilon > 0$, where ϵ represents a lower bound of the area reduction at each inductive step.

We claim that, due to our repair semantics, this lower bound ϵ depends on the initial instance D , and not on n . In order to prove this, let us first remark that an admissible region is fully determined by its boundaries. Now we prove that the regions in any accessible instance depend on the regions in the original database instance, or, more precisely, by the boundaries delimiting those regions. We prove it by induction on the number of inductive steps of the definition of accessible instances.

First, we prove that it works for the first repair transformation on the original database instance. Let $g'_1 = Tr^T(g_1, g_2)$ be the first transformation applied on region g_1 to create the accessible instance D_1 from the original database instance D . For $T \in \mathcal{T} \setminus \{TO, IT\}$, and following the definitions of admissible transformations in Table 1, the geometry g'_1 is either g_\emptyset or a region whose boundary is formed by parts of the boundaries that limit regions g_1 and g_2 (see example of overlapping regions in Figure ??). For $T \in \{TO, IT\}$, g'_1 is formed by parts of the boundary of region g_1 and the boundary created by buffering d around g_2 . So, in this case, the boundary of g'_1 depends exclusively on the boundaries of g_1 and g_2 , and of the constant d .

Let assume that the geometries in an accessible instance D_n obtained after n inductive steps are g_\emptyset or regions whose boundaries depend on the original instance. In the $n + 1$ inductive step, another transformation $g'_i = Tr^{T'}(g_i, g_j)$ is applied. Following the definition of admissible transformations, g'_i becomes the geometry g_\emptyset or a region whose boundary is formed by part of the boundary of g_i and part of the boundary of g_j , as in the first inductive step. Thus, g'_i also depends on the original database instance. This establishes our claim.

Now, by the Archimedean property of real numbers, there is a number M such that $A_0 - M\epsilon < 0$. Thus, after a finite number of iterations (i.e. applications of conflict-solving operators), we reach a consistent instance or an instance with area 0, i.e. all of whose geometries are empty, which would be consistent too.

Notice that the number M provides an upper bound on the number of times we can apply operators to produce a repair. At each point we have a finite number of choices. So, the overall number of o-repairs that can be produced is finite. \square

The following example shows that, even when applying admissible transformations, there may be exponentially many minimal repairs in the size of the database, a phenomenon already observed with relational repairs with respect to functional dependencies [6].

Example 8 Consider the schema in Example 2, and the SIC (3). The database instance contains n spatial tuples, as shown in Figure 9. There are $n-1$ overlappings and n overlapping geometries.




Figure 10: Exponential number of repairs

In order to solve each of those overlaps, we have the options of shrinking either one of the two regions involved. We have 2^{n-1} possible minimal repairs. \square

The following remark is important when estimating the data complexity of repairs, because, in this case, data complexity does not only depend on the number of tuples, but also on the size of geometric representations.

Remark 2 Transformation operators that make geometries empty reduce the size of geometric representations. Any other admissible transformation operator $tr^T(g_1, g_2)$ shrinks g_1 , and uses g_1 and g_2 to define the new boundary of g_1 . Thus, we are using, in a simple manner, the original geometric representation (e.g. points in the boundaries of the original geometries) to define a new geometry. It is clearly the case that there is a polynomial upper bound on the size of the representation of a new geometry in an o-repair in terms of the size of the original database, including representations of geometric regions. \square

4 Consistent Query Answers

We can use the concept of minimal repairs as an auxiliary concept to define, and possibly compute, consistent answers to a relevant class of queries in $\mathcal{L}(\Sigma)$.

A general conjunctive query is of the form:

$$\mathcal{Q}(\bar{v}) : \exists \bar{y} (R_1(\bar{x}_1; s_1) \wedge \cdots \wedge R_n(\bar{x}_n; s_n) \wedge \varphi), \quad (5)$$

where $\bar{v} = (\bigcup_i (\bar{x}_i \cup \{s_i\})) \setminus \bar{y}$ are the free variables, and φ is a conjunction of built-in atoms over thematic attributes or over spatial attributes that involve topological predicates in \mathcal{T} and geometric operators in \mathcal{O} . We also add a *safety condition*, requiring that variables in φ also appear in some of the R_i . For example, the following is a conjunctive query:

$$\mathcal{Q}(x, y; s) : \exists s_1 s_2 (R(x; s_1) \wedge R(y; s_2) \wedge Intersects(s_1, s_2) \wedge x \neq y \wedge s = Difference(s_1, s_2)).$$

We will consider the simpler but common and relevant class of conjunctive queries that are *operator free*, i.e., conjunctive queries of the form (5) where φ does not contain geometric operators. We will also study in more detail two particular classes of conjunctive queries:

(a) *Spatial range queries* are of the form

$$\mathcal{Q}(\bar{u}; s) : \exists \bar{z}(R(\bar{x}; s) \wedge T(s, w)), \quad (6)$$

with w a spatial constant, and $\bar{z} \subseteq \bar{x}$. This is a “query window”, and its free variables are those in $\bar{u} = (\bar{x} \setminus \bar{z})$ or in $\{s\}$.

(b) *Spatial join queries* are of the form

$$\mathcal{Q}(\bar{u}; s_1, s_2) : \exists \bar{z}(R_1(\bar{x}_1; s_1) \wedge R_2(\bar{x}_2; s_2) \wedge T(s_1, s_2)), \quad (7)$$

with $T \in \mathcal{T}$, and $\bar{z} \subseteq \bar{x}_1 \cup \bar{x}_2$. The free variables are those in $\bar{u} = ((\bar{x}_1 \cup \bar{x}_2) \setminus \bar{z})$ or in $\{s_1, s_2\}$.

We call *basic conjunctive queries* to queries of the form (6) or (7) with $T \in \{IIintersects, Intersects\}$.

Remark 3 Notice that for these two classes of queries we project on all the geometric attributes. We will also assume that the free variables correspond to a set of attributes of R with its key of the form (1). More precisely, for range queries, the attributes associated with \bar{u} contain the key of R . For join queries, $\bar{u} \cap \bar{x}_1$ and $\bar{u} \cap \bar{x}_2$ contain the key for relations R_1, R_2 , respectively. This is a common situation in spatial databases, where a geometry is retrieved together with its key values. \square

Given a query $\mathcal{Q}(\bar{x}; \bar{s})$, with free thematic variables \bar{x} and free geometric variables \bar{s} , a sequence of the thematic/spatial constants $\langle \bar{c}; \bar{g} \rangle$ is an answer to the query in instance D if and only if $D \models \mathcal{Q}(\bar{c}; \bar{g})$, that is the query \mathcal{Q} becomes true in D as a formula when its free variables \bar{x}, \bar{s} are replaced by the constants in \bar{c}, \bar{g} , respectively. We denote with $\mathcal{Q}(D)$ the set of answers to \mathcal{Q} in instance D .

Example 9 Figure 10 shows an instance for the schema $\mathcal{R} = \{LandP(idl; geometry), Building(idb; geometry)\}$. Here, idl, idb are keys for their relations. Dark rectangles represent buildings, and white rectangles represent land parcels. The queries \mathcal{Q}_1 and \mathcal{Q}_2 below are a range and a join query, respectively. For the former, the spatial constant is the spatial window shown in Figure 10, namely the (closed) polygon obtained by joining the four points in order indicated in the query.

$$\begin{aligned} \mathcal{Q}_1(idb; g) & : Building(idb; g) \wedge \\ & Intersects(g, ([x_1, y_1], [x_2, y_1], [x_2, y_2], [x_1, y_2], [x_1, y_1])). \\ \mathcal{Q}_2(idl, idl'; g, g') & : LandP(idb; g) \wedge LandP(idb'; g') \wedge Touches(g, g'). \end{aligned}$$

The answer to \mathcal{Q}_1 is $\langle idb_2; g_5 \rangle$. The answers to \mathcal{Q}_2 are: $\{\langle idl_1, idl_2; g_1, g_2 \rangle, \langle idl_2, idl_3; g_2, g_3 \rangle, \langle idl_1, idl_3; g_1, g_3 \rangle, \langle idl_2, idl_1; g_2, g_1 \rangle, \langle idl_3, idl_2; g_3, g_2 \rangle, \langle idl_3, idl_1; g_3, g_1 \rangle\}$. \square

Now we define the notion of consistent answer to a conjunctive query.

Definition 9 Consider an instance D , a set Ψ of SICs, and a conjunctive query $\mathcal{Q}(\bar{x}; \bar{s})$. A tuple of thematic/geometric constants $\langle c_1, \dots, c_m; g_1, \dots, g_l \rangle$ is a *consistent answer* to \mathcal{Q} with respect to Ψ if: (a) For every $D' \in Rep(D, \Psi)$, there exist g'_1, \dots, g'_l such that $D' \models \mathcal{Q}(c_1, \dots, c_m; g'_1, \dots, g'_l)$. (b) g_i is the intersection over all regions g'_i that satisfy (a) and are correlated to the same tuple in D .⁹ $Con(\mathcal{Q}, D, \Psi)$ denotes the set of consistent answers to \mathcal{Q} in instance D with respect to Ψ . \square

⁹Via the correlation function f , cf. Definition 3.




Figure 11: Example of a range query




Figure 12: Consistent answers

Since Q is operator free, the regions g'_i appear in relations of the repairs, and then f^{-1} can be applied. However, due to the intersection of geometries, the geometries in a consistent answer may not belong to the original instance or to any of its repairs.

In contrast to the definition of consistent answer to a relational query [2], where a consistent answer is an answer in every repair, here we have an aggregation of query answers via the geometric intersection and grouped-by thematic tuples. This definition is similar in spirit to consistent answers to aggregate relational queries with group-by [3, 9, 17].

This definition of consistent answer allows us to obtain more significative answers than in the relational case, because when shrinking geometries, we cannot expect to have, for a fixed tuple of thematic attribute values, the same geometry in every repair. If we did not use the intersection of geometries, we might lose or not have consistent answers due to the lack of geometries in common among repairs.

Example 10 (example 6 cont.) Consider the spatial range query

$$Q(idl; geometry) : \exists name owner (LandP(idl, name, owner; geometry) \wedge Intersects(geometry, ([x_1, y_1], [x_2, y_1], [x_2, y_2], [x_1, y_2], [x_1, y_1])),$$

which is expressed in the SQL language as:

```

SELECT  idl, geometry
FROM    LandP
WHERE   Intersects(geometry, ([x1, y1], [x2, y1], [x2, y2], [x1, y2], [x1, y1])).

```

Now, consider the two minimal repairs in Figure 6. In them, objects idl_1 and idl_3 do not change geometries, whereas object idl_2 does, from g_2 to g'_2, g''_2 , resp. (cf. Figure 6(a), (b), resp.).

From the first repair we get the following (usual) answers to the query: $\langle idl_1; g_1 \rangle, \langle idl_2; g'_2 \rangle, \langle idl_3; g_3 \rangle$. From the second repair, we obtain $\langle idl_1; g_1 \rangle, \langle idl_2; g''_2 \rangle, \langle idl_3; g_3 \rangle$. The consistent answers are the tuples shown in Figure 11, where the answers obtained in the repairs are grouped by an idl in common, and the associated geometries are intersected. In this figure, the geometry with thicker lines corresponds to the intersection of geometries obtained from different repairs.

From a practical point of view, the consistent query answer could include additional information about the degree in which geometries differ from their corresponding original geometries. For example, for the




Figure 13: The core of an instance

answer $\langle idl_2; g_2' \rangle$, an additional information could be the relative difference between areas g_2 and g_2' , which is calculated by $\delta(g_2, g_2')/area(g_2)$. \square

5 Core-Based CQA

The definition of consistent query answer relies on the auxiliary notion of minimal repair. However, since we may have a large number of repairs, computing consistent answers by computing, materializing, and finally querying all the repairs must be avoided whenever there are more efficient mechanisms at hand. Along these lines, in this section we present a methodology for computing consistent query answers to a subclass of conjunctive queries with respect to certain kind of SICs. It works in polynomial time (in data complexity), and does not require the explicit computation of the database repairs.

We start by defining the *core*, which is a single database instance associated with the class of repairs. We will use the core to consistently answer a subclass of conjunctive queries. Intuitively, the core is the “geometric intersection” of the repairs, which is obtained by intersecting the geometries in the different repair instances that correlate to the same thematic tuple.

Definition 10 For an instance D and a set Ψ of SICs, the *core* of D is the instance D^* given by $D^* := \{R(\bar{a}; g^*) \mid R \in \mathcal{R}, \text{ there is } R(\bar{a}; g) \in D \text{ and } g^* = \bigcap \{g' \mid R(\bar{a}; g') \in D' \text{ for some } D' \in Rep(D, \Psi) \text{ and } R(\bar{a}; g') = f_{D'}(R(\bar{a}; g))\}\}$. Here, $f_{D'}$ is the correlation function for D' .¹⁰ \square

Sometimes we will refer to D^* by $\bigcap^g Rep(D, \Psi)$. However, it cannot be understood as the set-theoretic intersection of the repairs of D . Rather it is a form of geometric intersection of geometries belonging to different repairs and grouped by common thematic attributes. It is important to remark that the keys of relations remain in the repairs, and therefore they appear in the core of a dimension instance.

Example 11 Figure 12 shows the *core* of the database instance in Figure 5 considering the repairs in Figure 6. Here, g_2^* results from the geometric intersection of geometries g_2' and g_2'' of the minimal repairs in Figure 6. Similarly, g_3^* is g_3' , because the latter is shared by both minimal repairs in Figure 6. Geometry g_4 becomes g_\emptyset in the core. All other geometries in the core are unchanged with respect to geometries in the original database instance. \square

Notice the resemblance between the definitions of consistent answer and the core. Actually, it is easy to see that $D^* = \bigcup_{R \in \mathcal{R}} Con(Q_R, D, \Psi)$, where the query $Q_R(\bar{x}; s) : R(\bar{x}; s)$ asks for the tuples in relation R .

The *core* is defined as the geometric intersection of all database repairs. However, as we will show, for a subset of SICs we can actually determine the *core* without computing these repairs. This is possible for

¹⁰Here, \bigcap is a set-theoretic intersection of geometries.

SICs of the form:

$$\forall \bar{x}_1 \bar{x}_2 s_1 s_2 \neg (R(\bar{x}_1; s_1) \wedge R(\bar{x}_2; s_2) \wedge \bar{x}'_1 \neq \bar{x}'_2 \wedge T(s_1, s_2)), \quad (8)$$

where $T \in \{IIntersects, Intersects, Equals\}$, $\bar{x}'_1 \subseteq \bar{x}_1$, $\bar{x}'_2 \subseteq \bar{x}_2$, and both \bar{x}'_1 and \bar{x}'_2 are the key of R . In these SICs there are two occurrences of the same database predicate in the same SIC. The following example illustrates this class of SICs.

Example 12 For the schema $\mathcal{R} = \{County(idc, name; geometry), Lake(idl; geometry)\}$, with idc the key of *County* and idl the key of *Lake*, the following SICs are of the form (8):

$$\bar{\forall} \neg (County(idc_1, n_1; s_1) \wedge County(idc_2, n_2; s_2) \wedge idc_1 \neq idc_2 \wedge IIntersects(s_1, s_2)). \quad (9)$$

$$\bar{\forall} \neg (Lake(idl_1; s_1) \wedge Lake(idl_2; s_2) \wedge idl_1 \neq idl_2 \wedge Intersects(s_1, s_2)). \quad (10)$$

□

Remark 4 This subset of SICs has the following properties, which will be useful when trying to compute the repairs and the core:

- (i) Two SICs of the form (8) over the same database predicate are redundant due to the semantic interrelation of the topological predicates *IIntersects*, *Intersects*, and *Equals*: only the constraint that contains the weakest topological predicate has to be considered. For example, *Intersects* is weaker than *IIntersects*, and *IIntersects* is weaker than *Equals*.
- (ii) Conflicts between tuples with respect to SICs of the form (8) are determined by the intersection of their geometries. The conflict between two tuples $R(\bar{a}_1; g_1)$ and $R(\bar{a}_2; g_2)$ is solved by applying a single admissible transformation operator $tr^T(g_1, g_2)$ (or $tr^{T^c}(g_2, g_1)$) that modifies g_1 (or g_2), and makes $T(g_1, g_2)$ (and $T^c(g_2, g_1)$) false.
- (iii) Solving conflicts with respect to a SIC of the form (8) is independent from solving a conflict with respect to another SIC of form (8) over a different database predicate.
- (iv) Solving a conflict between two tuples with respect to a SIC of the form (8) does not introduce new conflicts. This is due to the definition of admissible transformations and the monotonicity property of predicates *IIntersects* and *Intersects*, which prevent a shrunk geometry (or even an empty geometry) from participating in a new conflict with an existing geometry in the database (cf. Example 13).
- (v) For any two geometries g_1 and g_2 in conflict with respect to a SIC of the form (8), there always exist two repairs, one with the shrunk version of g_1 , and another with the shrunk version of g_2 . This guarantees that there exists a minimal repair that contains a minimum version of a geometry whose its geometric intersections with original geometries in conflict have been eliminated (cf. Lemma 2). As a consequence, the core can be computed by taking from a geometry all its intersections with other geometries in conflict, disregarding the order in which these intersections are eliminated.

This property is not guaranteed for other kinds of SICs. For instance, consider Example 6 with the instance in Figure 5 and its corresponding repairs in Figure 6. Although g_6 was originally in conflict with respect to g_2 , there is no minimal repair where geometry g_6 is shrunk. □

We illustrate some of these properties with the following example.




Figure 14: An inconsistent database with SICs of the form (8)

Example 13 (example 12 cont.) Consider the inconsistent instance in Figure 13. In it, counties with geometries g_1 , g_2 and g_3 are inconsistent with respect to SIC (9), because they internally intersect. Also, county g_3 internally intersects with geometry g_4 . Lakes with geometries g_6 and g_7 violate SIC (10), because they intersect (actually they touch).

Conflicts with respect to SICs (9) and (10) can be solved in an independent way, since they do not share predicates (cf. Remark 3(iii)). To obtain a repair, consider first SIC (9) and the conflict between g_2 and g_3 , which is solved by applying $tr^{II}(g_2, g_3)$ or $tr^{II}(g_3, g_2)$. Any of these alternative transformations do not produce geometries that could be in conflict with other geometries unless they were originally in conflict (cf. Remark 3(iv)). For instance, if we apply $tr^{II}(g_3, g_2)$ we obtain a new geometry g'_3 that will be in conflict with geometries g_1 and g_4 . These conflicts are not new, since g_3 was originally in conflict with these two geometries. Even more, by shrinking g_2 or g_3 , none of the modified geometries could be in conflict with g_5 . In addition, although by making $g'_3 = tr^{II}(g_3, g_2)$ we also solve the conflict between g_1 and g_3 , this is only accomplished due to the fact the conflicting part of g_3 and g_1 has been already eliminated from g'_3 (cf. Remark 3(v)).

Figure 14 shows the sixteen possible minimal repairs that are obtained by considering the eight possible ways in which conflicts with respect to SIC (9) are solved, in combination with the two possible ways in which conflicts with respect to SIC (10) are solved. In this figure thick boundaries represent geometries that have changed. Notice that in this figure we only show g'_{2_1} and not g'_{2_2} , since the later corresponds to the empty geometry which is then omitted in the corresponding repairs. The core for this database instance is shown in Figure 15. \square

It is possible to use a tree to represent all the versions that a geometry g may take in the repairs. The root of that tree is the original geometry g , the leaves are all the possible versions of the g in the minimal repairs. The internal nodes represent partial transformations applied to g as different conflicts in which g participates are solved. For illustration, Figure 16 shows the tree that represents the possible different versions of g_3 in the minimal repairs for the inconsistent instance in Figure 13. Notice that a leaf in this tree represents a version of g_3 in a repair, which is not necessarily a minimum geometry. For instance, in Figure 16 the minimum version of g_3 is g'_{3_1} . For all other non-minimum versions of g_3 in the leaves, conflicting areas are taken from other geometries. For example, geometry g'_{3_6} results by keeping g_3 as originally and shrinking geometries g_1 , g_2 and g_4 .

The following lemma establishes that when a geometry g is involved in conflicts of SICs of the form (8), there exists a version of g in the repairs that is minimum with respect to set-theoretic (geometric) inclusion. This result is useful to show that the minimum version of g is the one that will be in the core.

We need to introduce the set $\mathcal{G}_{R, \Psi}(\bar{a}, g)$ that contains, for a given tuple $R(\bar{a}; g)$ in a database instance D , all the possible versions of geometry g in the minimal repairs of D .

Definition 11 Let D a database instance, a set Ψ of SICs of the form (8) and a fixed tuple $R(\bar{a}; g) \in D$.



Figure 15: The sixteen possible repairs of Example 13

County*		
<i>idl</i>	<i>name</i>	<i>geometry</i>
<i>idc</i> ₁	<i>n</i> ₁	<i>g</i> ' _{1,1}
<i>idc</i> ₂	<i>n</i> ₂	<i>g</i> _∅
<i>idc</i> ₃	<i>n</i> ₃	<i>g</i> ' _{3,1}
<i>idc</i> ₄	<i>n</i> ₄	<i>g</i> ' _{4,1}
<i>idc</i> ₄	<i>n</i> ₅	<i>g</i> ' _{5,1}

Lake*	
<i>idl</i>	<i>geometry</i>
<i>idl</i> ₁	<i>g</i> ' _{6,1}
<i>idl</i> ₂	<i>g</i> ' _{7,1}




Figure 16: The core of Example 13




Figure 17: A tree-based representation of derived geometries from g_3 in some of the possible minimal repairs (thick boundaries represent geometries that have changed)

Then, $\mathcal{G}_{R,\Psi}(\bar{a}, g) = \{g' | R(\bar{a}; g') \in D', D' \in Rep(D, \Psi)\}$, $f^{-1}(R(\bar{a}; g')) = R(\bar{a}; g)$. □

Lemma 2 The set of geometries $\mathcal{G}_{R,\Psi}(\bar{a}, g)$ has a minimum element g_{min} under set-theoretic inclusion.

Proof: By properties of SICs of the form (8), for each conflict in which $R(\bar{a}; g) \in D$ participates, we can

or cannot shrink g . This leads to a combination of possible transformations over geometry g that can be represented in a binary tree as shown in Figure 16. So, we have a non-empty set of geometries $\mathcal{G}_{R,\Psi}(\bar{a}, g)$.

In this tree, we always have a path from the root to a leaf in which the geometry is always shrunk; that is, all conflicting areas are eliminated from g . The leaf geometry in this path (repair) is the minimum geometry g_{min} . \square

Corollary 1 Consider a database instance D , a set Ψ of SICs of the form (8), and a fixed tuple $R(\bar{a}; g) \in D$. For g_{min} , the minimum geometry in $\mathcal{G}_{R,\Psi}(\bar{a}, g)$, it holds $R(\bar{a}; g_{min}) \in D^*$.

Proof: Direct from Lemma 2 and the definition of the core as a geometric intersection. \square

5.1 Properties of the Core

In this section we establish that for the set of SICs of the form (8), and basic conjunctive queries, it is possible to compute consistent answers on the basis of the core of an inconsistent instance, avoiding the computation of queries in every minimal repair. This is established in Theorems 1 and 2, respectively.

Theorem 1 For an instance D , a set Ψ of SICs of the form (8), and a basic spatial range query $\mathcal{Q}(\bar{u}; s)$, it holds $\langle \bar{a}; g \rangle \in \text{Con}(\mathcal{Q}, D, \Psi)$ if and only if $\langle \bar{a}; g \rangle \in \mathcal{Q}(D^*)$.

Proof: The projection of range queries always includes the key of the relation in the result. Thus, if $\langle \bar{a}, g \rangle \in \text{Con}(\mathcal{Q}, D, \Psi)$, then for every $D' \in \text{Rep}(D, \Psi)$, there exists $R(\bar{b}; g')$, such that $\bar{a} \subseteq \bar{b}$, $f^{-1}(R(\bar{b}; g')) = R(\bar{b}; g)$ and $R(\bar{b}; g) \in D$, where $T(g', w)$ is true for the spatial constant w of the range query and $g = \bigcap g'$ with the intersection ranging over all g' .

By Lemma 2, there exists tuple $R(\bar{b}; g_{min}) \in D' \in \text{Rep}(D, \Psi)$ with $g_{min} \in \mathcal{G}_{R,\Psi}(\bar{b}; g)$. If $\langle \bar{a}; g \rangle \in \text{Con}(\mathcal{Q}, D, \Psi)$, with $\bar{a} \subseteq \bar{b}$, $g = \bigcap g' = g_{min}$. Also, it must happen that $\langle \bar{a}; g_{min} \rangle \in \mathcal{Q}(D')$. Then by Corollary 1, $\langle \bar{a}; g_{min} \rangle \in \mathcal{Q}(D^*)$, and therefore $\langle \bar{a}; g \rangle \in \mathcal{Q}(D^*)$.

In the other direction, if $\langle \bar{a}, g^* \rangle \in \mathcal{Q}(D^*)$ (with $D^* = \bigcap^g \text{Rep}(D, \Psi)$), then there exists a tuple $R(\bar{b}; g^*) \in D^*$, with $\bar{a} \subseteq \bar{b}$ and $g^* \neq g_{\emptyset}$. By the monotonicity of $T \in \{\text{Intersects}, \text{IIntersects}\}$, if $T(g^*, w)$ is true, then for all geometries g' in $R(\bar{b}; g^*) \in \text{Rep}(D, \Psi)$, with $g_{min} \subseteq g'$, $g' \neq g_{\emptyset}$, $T(g', w)$ is also true. Then, by Lemma 2 and Corollary 1, $g^* = \bigcap g' = g_{min}$ and $\langle \bar{a}, g^* \rangle \in \text{Con}(\mathcal{Q}, D, \Psi)$. \square

A similar result can be obtained for basic join queries, i.e., queries that consider two database predicates (not necessarily different). Notice that for a SIC φ of the form (8) with a database predicate R and a basic join query of the form (7) with $R = R_1 = R_2$, the consistent answers do not contain information from tuples that were originally in conflict. This is because by solving conflicts with respect to φ , all possible intersections between tuples in R will be eliminated (a basic join query asks for geometries that intersect).

The following example illustrates how to compute consistent answers to basic join queries. This example will also illustrate the proof of Theorem 2.

Example 14 (example 13 cont.) Consider the following basic join query posed to the instance D in Example 13. It is asking for the identifiers and geometries of counties and lakes that internally intersect.

$$\mathcal{Q}(idc, idl; g_1, g_2) : \exists n(\text{County}(idc, n; g_1) \wedge \text{Lake}(idl; g_2) \wedge \text{IIntersects}(g_1, g_2)).$$

The consistent answer to this query is $\langle idc_3, idl_2; g'_{3_1}, g'_{7_1} \rangle$. Without using the core, this answer is obtained by intersecting all answers obtained from every possible minimal repair. The geometries in the repairs of D with respect to Ψ (SICs (9) and (10)) can be partitioned into the following sets: $\mathcal{G}_{County, \Psi}(idc_1, n_1; g_1) = \{g'_{1_1}, g'_{1_2}\}$, $\mathcal{G}_{County, \Psi}(idc_2, n_2; g_2) = \{g_{\emptyset}, g'_{2_1}, g'_{2_2}\}$, $\mathcal{G}_{County, \Psi}(idc_3, n_3; g_3) = \{g'_{3_1}, g'_{3_2}, g'_{3_3}, g'_{3_4}, g'_{3_5}, g'_{3_6}\}$, $\mathcal{G}_{County, \Psi}(idc_4, n_4; g_4) = \{g'_{4_1}, g'_{4_2}\}$, $\mathcal{G}_{County, \Psi}(idc_5, n_5; g_5) = \{g'_{5_1}\}$, $\mathcal{G}_{Lake, \Psi}(idl_6; g_6) = \{g'_{6_1}, g'_{6_2}\}$, $\mathcal{G}_{Lake, \Psi}(idl_7; g_7) = \{g'_{7_1}, g'_{7_2}\}$. The minimum geometries in these seven sets are g'_{1_1} , g_{\emptyset} (corresponding to the update of geometry g'_2), g'_{3_1} , g'_{4_1} , g'_{5_1} , g'_{6_1} , and g'_{7_1} , respectively.

Also, for the database predicates *County* and *Lake*, there are two sets containing the possible extensions of them in the repairs: $\{County(D') | D' \in Rep(D, \Psi)\}$, containing the eight versions of counties (first eight versions of counties in Figure 14); and $\{Lake(D') | D' \in Rep(D, \Psi)\}$, with the two instances of lakes (one with geometries g'_{6_2} and g'_{7_1} , and the other with geometries g'_{6_1} and g'_{7_2} in Figure 14). Note that the possible minimal repairs contain combinations of geometries in sets $\mathcal{G}_{County, \Psi}(idc, n; g)$ and $\mathcal{G}_{Lake, \Psi}(idl; g)$. In particular, there exists a repair that combines the minimum geometries g'_{3_1} and g'_{6_1} , and another repair that combines g'_{3_1} and g'_{7_1} .

If the topological predicate in the basic join query is satisfied by the combination of two minimum geometries, then other versions of these geometries in other repairs (which geometrically include the minimum geometries) will also satisfy it. In this example, g'_{3_1} and g'_{7_1} intersect and, by the monotonicity property of predicate *IIntersects*, all other versions of g_3 and g_7 in other repairs also intersect. As result, $\langle idc_3, idl_2; g'_{3_1}, g'_{7_1} \rangle$ is an answer to the query. Finally, by Corollary 1, g'_{3_1} and g'_{7_1} are in the core of the database instance and, therefore, $\langle idc_3, idl_2; g'_{3_1}, g'_{7_1} \rangle$ is also an answer to the query over the core. \square

Theorem 2 For an instance D , a set Ψ of SICs of the form (8), and a basic spatial join query $\mathcal{Q}(\bar{x}_1, \bar{x}_2; s_1, s_2)$, it holds $\langle \bar{a}_1, \bar{a}_2, g_1, g_2 \rangle \in Con(\mathcal{Q}, D, \Psi)$ if and only if $\langle \bar{a}_1, \bar{a}_2, g_1, g_2 \rangle \in \mathcal{Q}(D^*)$.

Proof: The projection of join queries also includes keys. Thus, if $\langle \bar{a}_1, \bar{a}_2; g_1, g_2 \rangle \in Con(\mathcal{Q}, D, \Psi)$, then we have tuples $R_1(\bar{b}_1; g'_1) \in D'$, $R_2(\bar{b}_2; g'_2) \in D'$, for every $D' \in Rep(D, \Psi)$ with $\bar{a}_1 \subseteq \bar{b}_1$, $\bar{a}_2 \subseteq \bar{b}_2$, and $T(g'_1, g'_2)$ true for T in \mathcal{Q} . Thus, g_1 is the intersection of all those g'_1 , and g_2 is the intersection of all those g'_2 .

First, note that if $R_1 = R_2$, only tuples that were not originally in conflict may be in the answer. These tuples will be trivially in the core, because no geometric transformations over their geometries are applied. Thus, their geometries will be in the answer, if and only if, they satisfy the topological predicate in the query.

By the property (iii) of SICs of the form (8) (cf. Remark 3), solving conflicts on two different database predicates R_1 and R_2 are independent. Let us assume that $\{R_1(D') | D' \in Rep(D, \Psi)\}$ and $\{R_2(D') | D' \in Rep(D, \Psi)\}$ are the different extensions of predicates R_1 and R_2 in all possible minimal repairs. Then, $Rep(D, \Psi)$ contains database instances that result from the combination of these two sets. Consequently, and using Lemma 2, for two given \bar{b}_1 and \bar{b}_2 , there exists a repair $D' \in Rep(D, \Psi)$ such that $R_1(\bar{b}_1; g'_{1_{min}}) \in D'$ and $R_2(\bar{b}_2; g'_{2_{min}}) \in D'$, where $g'_{1_{min}}$ is minimum in $\mathcal{G}_{R_1, \Psi}(\bar{b}_1, g_1)$ and $g'_{2_{min}}$ is minimum in $\mathcal{G}_{R_2, \Psi}(\bar{b}_2, g_2)$.

We now prove that if $\langle \bar{a}_1, \bar{a}_2; g_1, g_2 \rangle \in Con(\mathcal{Q}, D, \Psi)$, then $\langle \bar{a}_1, \bar{a}_2; g_1, g_2 \rangle \in \mathcal{Q}(D^*)$. By definition of consistent answer, if $\langle \bar{a}_1, \bar{a}_2; g_1, g_2 \rangle \in Con(\mathcal{Q}, D, \Psi)$, then $\langle \bar{a}_1, \bar{a}_2, g'_{1_{min}}, g'_{2_{min}} \rangle \in \mathcal{Q}(D')$. By Corollary 1, $\langle \bar{a}_1, \bar{a}_2; g_1, g_2 \rangle \in \mathcal{Q}(D^*)$, with $g_1 = g'_{1_{min}}$ and $g_2 = g'_{2_{min}}$.

In the other direction, if $\langle \bar{a}_1, \bar{a}_2, g_1^*, g_2^* \rangle \in \mathcal{Q}(D^*)$, then $\langle \bar{a}_1, \bar{a}_2, g_1^*, g_2^* \rangle \in Con(\mathcal{Q}, D, \Psi)$. By Corollary 1, $g_1^* = g'_{1_{min}}$ and $g_2^* = g'_{2_{min}}$, and $R_1(\bar{b}_1; g'_{1_{min}}) \in D^*$ and $R_2(\bar{b}_2; g'_{2_{min}}) \in D^*$. Then, by monotonicity property of predicate $T \in \{Intersects, IIntersects\}$ in \mathcal{Q} , if $T(g'_1, g'_2)$ is true, it is also true for all $R_1(\bar{b}_1; g''_1) \in D''$ and in $R_2(\bar{b}_2; g''_2) \in D''$, with $D'' \in Rep(D, \Psi)$ and $g'_{1_{min}} \subseteq g''_1$ and $g'_{2_{min}} \subseteq g''_2$. Therefore, $\langle \bar{a}_1, \bar{a}_2, g_1^*, g_2^* \rangle \in Con(\mathcal{Q}, D, \Psi)$. \square




Figure 18: Core vs. consistent answers

The previous theorems tell us that we can obtain consistent answer to basic conjunctive queries by direct and usual query evaluation on the single instance D^* , the *core* of D . This does not hold for non-basic conjunctive queries as the following example shows.

Example 15 Consider a database instance with a database predicate R whose geometric attribute values are shown in Figure 17(a). This database instance is inconsistent with respect to a SIC that specifies that geometries cannot overlap. Let us now consider a range query of the form $\exists \bar{y}(R(\bar{x}; g) \wedge Touches(g, s))$, where s is a user defined spatial window, and $\bar{y} \subseteq \bar{x}$. Figure 17(b) shows the query over the intersection of all repairs (the *core*), obtaining geometries g_1^* and g_2^* , from where only g_1^* touches s . Figures 17(c) and (d) show the query over each repair, separately. The answer from the repair in (c) is g_1' , and repair (d) does not return an answer because none of the geometries in this repair touches s . Their intersection, therefore, is empty and differs from the answer obtained from the *core*. This difference is due to the fact that the query window s touches geometry g_1' in only one of the repairs. \square

5.2 Computing the Core

We now give a characterization of the core of a database instance for a set of SICs of the form (8), which is not explicitly based on the computation of minimal repairs. This equivalent and alternative characterization of the core allows us to compute the core without having to compute all the minimal repairs.

To simplify the notation, we introduce a logical formula that captures a conflict around a tuple of relation $R \in D$ and a SIC of the form (8) with topological predicate T :

$$\forall \bar{x}_1 \bar{x}_2 s_1 s_2 (ConfI_{D,R,T}(\bar{x}_1, s_1, \bar{x}_2, s_2) \equiv (R(\bar{x}_1; s_1) \wedge R(\bar{x}_2; s_2) \wedge \bar{x}_1 \neq \bar{x}_2 \wedge T(s_1, s_2))). \quad (11)$$

Definition 12 Let D be a database instance and Ψ a set of SICs of the form (8). For the *core* D^* of D with respect to Ψ , it holds $D^* = D_{Intersects}^* \cap D_{Intersects}^* \cap D_{Equal}^*$, where:

- (a) $D_{Intersects}^* = \{R(\bar{a}; Difference(g, t)) \mid R(\bar{a}; g) \in D, t = \bigcup \{g' \mid \text{for every } R(\bar{b}, g') \in D \text{ such that } D \models ConfI_{D,R,Intersects}(\bar{a}, g, \bar{b}, g')\}\}$, where \bigcup is the *geomUnion* operator that calculates the geometric union (spatial aggregation) of geometries.

- (b) $D_{Intersects}^* = \{R(\bar{a}; Difference(g, t)) \mid R(\bar{a}; g) \in D, t = \bigcup \{Buffer(g', d) \mid \text{for every } R(\bar{b}, g') \in D \text{ such that } D \models ConfI_{D,R,Intersects}(\bar{a}, g, \bar{b}, g')\}\}$.
- (c) $D_{Equal}^* = \{R(\bar{a}; g) \mid R(\bar{a}, g) \in D, \text{it does not exists } (R(\bar{b}, g) \in D, D \models ConfI_{D,R,Equal}(\bar{a}, g, \bar{b}, g))\}$.
□

Notice that t is the union of all the geometries that are in conflict with a given geometry g . It is obtained by using the spatial aggregation operator *geomUnion*.

Now, we give the specification of the cores: $D_{Intersects}^*$, $D_{IIntersects}^*$,¹¹ and D_{Equal}^* as views in a spatial SQL language.¹² In the following specification, we assume a database instance D with a relational predicate $R(id; geometry)$ and primary key id . The following specification shows that our methodologies could be implemented on top of current spatial database management systems. In particular, the definition of $D_{Intersects}^*$ uses a fixed value d that represents the minimum distance between geometries in the cartographic scale of the database instance. The intersection of these views makes D^* .

Table 3 shows three views that enables to compute the core of the database with a database predicate $R(idl; geometry)$.

Example 16 (example 10 cont.) The example considers only the relation *LandP* with primary key *idl* and the SIC (3) of Example 2. We want to consistently answer the query of Example 11, i.e., $\exists name \ owner (LandP(idl, name, owner; geometry) \wedge IIntersects(geometry, ([x_1, y_1], [x_2, y_1], [x_2, y_2], [x_1, y_2], [x_1, y_1])))$.

To answer this query, we generate a view of the *core* applying the definition in Table 3. That is, we eliminate from each geometry the union of conflicting regions with respect to each land parcel. In this case, the conflicting geometries for g_2 are g_3 and g_4 ; for geometry g_3 is g_2 ; and for geometry g_4 is g_2 . This is the definition of the core in SQL:

```
CREATE VIEW Core
AS (SELECT l1.idl AS idl, l1.name AS name, l1.owner AS owner,
difference(l1.geometry, geomunion(l2.geometry)) AS geometry
FROM LandP AS l1, LandP AS l2
WHERE l1.idl <> l2.idl AND Intersects(l1.geometry, l2.geometry) AND
NOT Touches(l1.geometry, l2.geometry)
GROUP BY l1.idl, l1.name, l1.owner, l1.geometry
UNION
SELECT l1.idl AS idl, l1.name AS name, l1.owner AS owner, l1.geometry AS geometry
FROM LandP AS l1
WHERE NOT EXISTS(SELECT l2.idl, l2.geometry
FROM LandP AS l2
WHERE l1.idl <> l2.idl AND Intersects(l1.geometry, l2.geometry) AND
NOT Touches(l1.geometry, l2.geometry)))
```

We now can pose the query to the core to compute the consistent answer to the original query:

¹¹In current SQL Language $IIntersects(g_1, g_2) = Intersects(g_1, g_2) \text{ AND NOT Touches}(g_1, g_2) = Overlaps(g_1, g_2) \text{ OR Within}(g_1, g_2) \text{ OR Contains}(g_1, g_2) \text{ OR Touches}(g_1, g_2)$.

¹²Optimizations to the SQL statements are possible by using materialized views and avoiding double computation of join operations.

$D_{Intersects}^*$	<pre> CREATE VIEW Core_Intersects AS (SELECT FROM WHERE GROUP BY UNION SELECT FROM WHERE </pre>	<pre> r1.id AS id, difference(r1.geometry, Buffer(geomunion(r2.geometry), d)) AS geometry R AS r1, R AS r2 r1.id <> r2.id AND Intersects(r1.geometry, r2.geometry) r1.id, r1.geometry r1.id AS id, r1.geometry AS geometry R AS r1 NOT EXISTS (SELECT r2.id, r2.geometry FROM R AS r2 WHERE r1.id <> r2.id AND Intersects(r1.geometry, r2.geometry))) </pre>
$D_{IIntersects}^*$	<pre> CREATE VIEW Core_IIntersects AS (SELECT FROM WHERE GROUP BY UNION SELECT FROM WHERE </pre>	<pre> r1.id AS id, difference(r1.geometry, geomunion(r2.geometry)) AS geometry R AS r1, R AS r2 r1.id <> r2.id AND Intersects(r1.geometry, r2.geometry) AND NOT Touches(r1.geometry, r2.geometry) r1.id, r1.geometry r1.id AS id, r1.geometry AS geometry R AS r1 NOT EXISTS (SELECT r2.id, r2.geometry FROM R AS r2 WHERE r1.id <> r2.id AND Intersects(r1.geometry, r2.geometry) AND NOT Touches(r1.geometry, r2.geometry))) </pre>
D_{Equal}^*	<pre> CREATE VIEW Core_Equal AS (SELECT FROM WHERE </pre>	<pre> r1.id AS id, r1.geometry AS geometry R AS r1 NOT EXISTS (SELECT r2.id, r2.geometry FROM R AS r2 WHERE r1.id <> r2.id AND Equals(r1.geometry, r2.geometry))) </pre>

Table 4: SQL statements to compute views for $D_{Intersects}^*$, $D_{IIntersects}^*$, and D_{Equal}^*

```

SELECT idl, name, owner, geometry
FROM Core
WHERE Intersects(geometry, ([x1, y1], [x2, y1], [x2, y2], [x1, y2], [x1, y1]))

```

(12)

The answer is shown in Figure 11. This query is a classic selection from the *Core* view. □

This core-based method allows us to compute consistent answers in polynomial (quadratic) time (in data complexity) in cases where there can be exponentially many repairs. In Example 8, where we have 2^{n-1} minimal repairs, we can apply the query Q over the *core*, and we only have to compute the difference of a geometry with respect to the union of all other geometries in conflict. This corresponds to a polynomial time algorithm of order polynomial with respect to the size of the database instance.

6 Experimental Evaluation

In this section we analyze the results of the experimental evaluation we have done of the core-based CQA using synthetic and real data sets. The experiment includes a scalability analysis that compares the cost of CQA with increasing numbers of conflicting tuples and increasing sizes of database instances. We compare these results with respect to the direct evaluation of basic conjunctive queries over the inconsistent database (i.e., ignoring inconsistencies). The latter reflects the additional cost of computing consistent answers against computing queries that ignore inconsistencies.

6.1 Experimental Setup

We create synthetic databases to control the size of the database instance and the number of conflicting tuples. We use a database schema consisting of a single predicate $R(id; geometry)$, where id is the numeric key and $geometry$ is a spatial attribute of type polygon. We create three sets of synthetic database instances by considering SICs of the form (8) with different topological predicates:

Set	SIC
Equals	$\forall \neg(R(\bar{x}_1; s_1) \wedge R(\bar{x}_2; s_2) \wedge \bar{x}_1 \neq \bar{x}_2 \wedge Equals(s_1, s_2))$
Intersects	$\forall \neg(R(\bar{x}_1; s_1) \wedge R(\bar{x}_2; s_2) \wedge \bar{x}_1 \neq \bar{x}_2 \wedge Intersects(s_1, s_2))$
IIntersects	$\forall \neg(R(\bar{x}_1; s_1) \wedge R(\bar{x}_2; s_2) \wedge \bar{x}_1 \neq \bar{x}_2 \wedge IIntersects(s_1, s_2))$

For each set we create five consistent instances including 5,000, 10,000, 20,000, 30,000, and 40,000 tuples of homogeneously distributed spatial objects whose geometries are rectangles (i.e., 5 points per geometric representation of rectangles). Then, we create inconsistent instances with respect to the corresponding SICs in each set with 5%, 10%, 20%, 30%, and 40% of tuples in conflict. For database instances with a SIC and topological *Equals*, we create inconsistencies by duplicating geometries in a percentage of geometries. For database instances with a SIC and topological *IIntersects*, we create inconsistencies by making geometries overlap. Finally, for database instances with a SIC and topological *Intersects*, we create inconsistencies by making a percentage of geometries to touch.

Due to the spatial distribution of rectangles in the sets, the cores for database instances with SICs using topological predicates in $\{Intersects, IIntersects\}$ have the same numbers of points in their geometric representations than their original instances. For the set of database instances with SICs using topological predicate *Equals*, the numbers of points in the geometric representations of their cores are less than in the original databases, because we eliminate geometries as we restore consistency. Thus, we are not introducing additional storage costs in our experiments.

To have a better understanding of the computational cost of CQA, we also evaluate the cost of CQA over real and free available data of administrative boundaries of Chile [1]. Chilean administrative boundaries have complex shapes with many islands, specially, in the South of Chile (e.g., a region can have 891 islands). For the real database, we have two predicates *Counties* and *Provinces*. Notice that, at the conceptual level, *Provinces* are aggregations of *Counties*. In this experiment, however, we have used the source data as it is, creating separated tables for *Counties* and *Provinces* with independent spatial attributes. For this real database, we consider SIC of the form: $\forall \neg(R(\bar{x}_1; s_1) \wedge R(\bar{x}_2; s_2) \wedge \bar{x}_1 \neq \bar{x}_2 \wedge IIntersects(s_1, s_2))$, with R being *Counties* or *Provinces*.

Table 4 summaries the data sets for the experimental evaluation. The percentage of inconsistency is calculated as the number of tuple in any conflict over the total number of tuples. The geometric representation size is calculated as the number of points in the boundaries of a region.

Source	Name	Tuples	Inconsistency (%)	Geometric representation size
Synthetic	Equals	5,000-40,000	5-40	25,000-200,000
	IIntersects	5,000-40,000	5-40	25,000-200,000
	Intersects	5,000-40,000	5-40	25,000-200,000
Real	Provinces	52	59	35,436
	Counties	307	12.7	72,009

Table 5: Data sets of the experimental evaluation

We measure the computational cost in terms of seconds needed to compute the SQL statement on a Quad Core Xeon X3220 of 2.4 GHz, 1066 MHz, and 4 GB in RAM. We use as spatial DBMS PostgreSQL 8.3.5 with PostGIS 1.3.5.

6.2 Experimental Results

Figure 18 shows the cost of the core computation for the different synthetic database instances. To make this experimental evaluation easier and faster, we used materialized views so that we computed only once the core and applied queries on this core’s view. However, we added the computational cost of the core to each individual query result to have a better understanding of the cost of applying CQA.

The time cost of computing the core for inconsistent databases with respect to a SIC with a topological predicate *Equals* decreases as the number of tuples in conflicts increases, since the core eliminates geometries in conflict and, therefore, these empty geometries are then ignored in geometric computations. The cost of computing the core is largely due to the spatial join given by the topological predicate of a SIC, which could decrease using more efficient algorithms and spatial indexing structures.




Figure 19: Time cost of the core computation for different SICs, different levels of inconsistency, and different sizes of databases instances

For the synthetic database instance, Figures 19 and 20 show the cost rate between computing a CQA with respect to simple range or join queries (with the spatial predicate *Intersects*) that ignore inconsistencies.

Range queries use a random query window created by a rectangle whose side is equivalent to 1% of the total length in each dimension. Notice that the time cost of computing a range query for a database instance with 10,000 was approximately 15 ms, which, in average, was 900 times less than computing a join query. These reference values exhibit linear and quadratic growth for range and join queries, respectively, as we consider increasing sizes of database instances. The computational cost of CQA to join queries include the computation of the core; however, this cost could be amortized if we use a materialized view of the core for computing more than one join query. In the time cost of CQA for range queries, we have optimized the computation by applying the core-computation over a subset of tuples previously selected by the query range. This optimization is not possible for join queries, since no spatial window can constrain the possible geometries in the answer.




Figure 20: Relative cost of CQA to range queries

The results indicate that CQA to a range query can cost 100 times the cost of a simple query. This is primarily due to the join computation of the core. Indeed, when comparing the CQA to a join query, we only duplicate the relative cost, and in the best case, keep the same cost. However, join queries have a significant larger computational cost. Notice that the computation cost for a CQA to range query is around 60s in the worst case (40,000 tuples). With exception of cases when the core contains empty geometries, the percentage of inconsistencies does not affect drastically the results.

We also evaluate the scalability of the CQA cost to range queries in function of the size of the query window (i.e., spatial window). In Figure 21 we show the relative CQA cost to range queries on a synthetic database instance with 10,000 tuples and range queries whose random spatial windows varied from 1% to 5% of the size in each dimension. The results indicate that the relative cost increases logarithmical as we increase the size of the query window. Also, only for database instances with a SIC and topological predicate *Equals*, the relative cost suffers some variation across different percentages of inconsistencies, primarily, due to the elimination of geometries in the database.

Finally, we applied the core-based computation of CQA to the real database instances in Table 4. Table 5 summaries the results obtained with these data, which were in agreement with the results obtained with the synthetic database instances. In this table, $\Delta Points$ represents the relative difference in the size of the geo-




Figure 21: Relative cost of CQA to join queries




Figure 22: Relative cost of CQA to range queries and different sizes of the query window (using a database instance with 10,000 tuples)

metric representation between the core and the original database. Notice that computing the core increased the geometric representation of *Provinces* up to 5.0%, which is bounded by the shape of geometries in conflict (i.e., the size of the original geometric representation). In the case of *Counties*, however, the size of the geometric representation of the core decreases down to -0.03% . Since the geometry of provinces should be the geometric aggregation of counties, we could expect to have a relationship between $\Delta Points$ for *Provinces* and *Counties*. However, the source data set uses independent geometries for *Provinces* and *Counties* and no comparison can be made.

Data	Δ Points	Core	Range		Join	
			Simple	CQA	Simple	CQA
Provinces	+5.0%	17.7	0.04	0.25	29.8	63.4
Counties	-0.03%	18.1	0.1	2.1	40.6	55.7

Table 6: CQA cost with real data (costs of core and queries in seconds)

7 Conclusions

We have formalized a repair semantics and consistency query answers for spatial databases with respect to SICs. The repair semantics is used as an auxiliary concept for handling inconsistency tolerance and computing consistent answers to spatial queries. It is based on updates that shrink geometries of objects, even at the point of deleting geometries for some exceptional cases, as for predicate *Disjoint*. Geometries are virtually updated applying admissible geometric operators, which are available in most spatial DBMSs.

By restricting ourselves to the application of the admissible transformations, we have a finite number of possibilities for making a pair of geometries consistent with respect to a SIC. However, there may still be exponentially many repairs for a given instance and set of SICs. With the purpose of avoiding to compute and query all repairs, we have identified cases of SICs and conjunctive (range and join) queries where the consistent answers can be obtained by posing a standard query to a single view of the original instance. This view is equivalent to the intersection of all possible minimal repairs, what we called the *core* of a database instance, which for a subset of SICs can be computed in polynomial time without determining each repair.

An experimental evaluation of the core-based computation of CQA reveals that answering range queries has a cost that varies drastically in function of the topological predicates in SICs and the number of tuples in the database instance, reaching up to 100 times the cost of a simple range query. This is mainly due to the spatial join involved in computing the core. For join queries, instead, the cost of CQA is the double of a simple join query. These results do not use optimizations with spatial indexing, which has been left for future work. Even more, they assume that we have to compute the core for each query, which could be optimized by using materialized views.

This work leaves many problems open. Most prominently, computability and complexity issues have to be explored. For example, some interesting decision problems are deciding if non trivial repairs (i.e., not obtained by cancellation of geometries) exist for an instance and a set of SICs, or deciding whether or not a particular instance is a repair of an inconsistent database instance. The complexity of deciding if a spatio-relational tuple is a consistent answer is also open. As in the relational case, we expect to find hard cases for all these problems. For them, it would be interesting to obtain lower complexity approximation algorithms.

We have considered only regions to represent spatial objects. A natural extension of this work would be to define a repair semantics for other spatial abstractions, such as polylines, points, networks, and so on. We would also like to explore not only denial SICs, but also other classes of semantic ICs, and other types of repair semantics that include solving conflicts with respect to a topological predicate *Disjoint*. This includes also the possibility of considering combinations of spatial with relational constraints, e.g. functional dependencies and referential ICs.

Acknowledgments

This project is partially funded by FONDECYT, Chile, grant number 1080138. Part of this research was done when L. Bertossi was invited to the Universidad de Concepcion and Universidad del Bío-Bío. Mónica

Caniupán received funding from FONDECYT, Chile, grant number 11070186. Leopoldo Bertossi has also been partially funded by an NSERC Discovery Grant (#315682).

References

- [1] ALBERS, C. Gisdata chile, administrative boundaries since 2007. <http://www.rulamahue.cl/mapoteca/catalogos/chile.html>, 2009.
- [2] ARENAS, M., BERTOSI, L., AND CHOMICKI, J. Consistent query answers in inconsistent databases. In *18th ACM Symposium on Principles of Database Systems PODS'99* (1999), ACM Press, pp. 68–79.
- [3] ARENAS, M., BERTOSI, L., AND CHOMICKI, J. Scalar aggregation in fd-inconsistent databases. In *Proc. International Conference on Database Theory* (2001), Springer LNCS 1973, pp. 39–53.
- [4] BERTOSI, L. Consistent query answering in databases. *ACM Sigmod Record* 35, 2 (2006), 68–76.
- [5] BERTOSI, L., AND BRAVO, L. Consistent query answers in virtual data integration systems. In *Inconsistency Tolerance* (2005), vol. 3300 of *Lecture Notes in Computer Science*, Springer, pp. 42–83.
- [6] BERTOSI, L., AND CHOMICKI, J. *Logics for Emerging Applications of Databases*. Springer, 2003, ch. Query Answering in Inconsistent Databases, pp. 43–83.
- [7] BERTOSI, L., HUNTER, A., AND SCHAUB, T., Eds. *Inconsistency Tolerance* (2005), vol. 3300 of *Lecture Notes in Computer Science*, Springer.
- [8] BORGES, K., LAENDER, A., AND DAVIS, C. Spatial integrity constraints in object oriented geographic data modeling. In *ACM International Symposium on Advances in GIS* (1999), ACM Press, pp. 1–6.
- [9] CANIUPAN, M. *Optimizing and Implementing Repair Programs for Consistent Query Answering in Databases*. PhD thesis, Carleton University, Department of Computer Science, 2007.
- [10] CHOMICKI, J. Consistent query answering: Five easy pieces. In *International Conference on Database Theory ICDT'07* (2007), vol. LNCS 4353, Springer-Verlag, pp. 1–17.
- [11] COCKCROFT, S. A taxonomy of spatial integrity constraints. *GeoInfomatica* 1, 4 (1997), 327–343.
- [12] DUCKHAM, M., LINGHAM, J., MASON, K. T., AND WORBOYS, M. F. Qualitative reasoning about consistency in geographic information. *Inf. Sci.* 176, 6 (2006), 601–627.
- [13] EGENHOFER, M., CLEMENTINE, E., AND FELICE, P. D. Evaluating inconsistency among multiple representations. In *Spatial Data Handling* (1995), pp. 901–920.
- [14] EGENHOFER, M., AND SHARMA, J. Assessing the consistency of complete and incomplete topological information. *Geographical Systems* 1 (1993), 47–68.
- [15] EGENHOFER, M. J., AND FRANZOSA, R. D. Point set topological relations. *International Journal of Geographical Information Systems* 5 (1991), 161–174.

- [16] FRANCONI, E., PALMA, A. L., LEONE, N., PERRI, S., AND SCARCELLO, F. Census data repair: a challenging application of disjunctive logic programming. In *LPAR (2001)*, R. Nieuwenhuis and A. Voronkov, Eds., vol. 2250 of *Lecture Notes in Computer Science*, Springer, pp. 561–578.
- [17] FUXMAN, A., FAZLI, E., AND MILLER, R. J. Conquer: Efficient management of inconsistent databases. In *SIGMOD Conference (2005)*, F. Özcan, Ed., ACM, pp. 155–166.
- [18] GÜTING, R. H. Graphdb: Modeling and querying graphs in databases. In *VLDB (1994)*, J. B. Bocca, M. Jarke, and C. Zaniolo, Eds., Morgan Kaufmann, pp. 297–308.
- [19] GÜTING, R. H., AND SCHNEIDER, M. Realm-based spatial data types: The rose algebra. *VLDB J.* 4, 2 (1995), 243–286.
- [20] HADZILACOS, T., AND TRYFONA, N. A model for expressing topological integrity constraints in geographic databases. In *Spatio-Temporal Reasoning (1992)*, A. U. Frank, I. Campari, and U. Formentini, Eds., vol. 639 of *Lecture Notes in Computer Science*, Springer, pp. 252–268.
- [21] KUIJPERS, B., PAREDAENS, J., AND DEN BUSSCHE, J. V. On topological elementary equivalence of spatial databases. In *ICDT (1997)*, F. N. Afrati and P. G. Kolaitis, Eds., vol. 1186 of *Lecture Notes in Computer Science*, Springer, pp. 432–446.
- [22] MÄS, S. Reasoning on spatial semantic integrity constraints. In *COSIT (2007)*, S. Winter, M. Duckham, L. Kulik, and B. Kuipers, Eds., vol. 4736 of *Lecture Notes in Computer Science*, Springer, pp. 285–302.
- [23] OPENGIS. Opengis simple features specification for sql. Tech. rep., Open GIS Consortium, 1999.
- [24] PAREDAENS, J., DEN BUSSCHE, J. V., AND GUCHT, D. V. Towards a theory of spatial database queries. In *PODS (1994)*, ACM Press, pp. 279–288.
- [25] PAREDAENS, J., AND KUIJPERS, B. Data models and query languages for spatial databases. *Data Knowl. Eng.* 25, 1-2 (1998), 29–53.
- [26] RANDELL, D. A., CUI, Z., AND COHN, A. G. A spatial logic based on regions and connection. In *KR (1992)*, pp. 165–176.
- [27] RODRÍGUEZ, M. A. Inconsistency issues in spatial databases. In *Inconsistency Tolerance (2005)*, vol. 3300 of *Lecture Notes in Computer Science*, Springer, pp. 237–269.
- [28] RODRÍGUEZ, M. A., BERTOSSI, L., AND CANIUPAN, M. An inconsistency tolerant approach to querying spatial databases. In *ACM-GIS (2008)*, W. G. Aref, M. Mokbel, and M. Schneider, Eds., ACM.
- [29] SERVIGE, S., PURICELLI, T., AND LAURINI, R. A methodology for spatial consistency improvement of geographic databases. *GeoInformatica* 4 (2000), 7–24.
- [30] THOMPSON, D., AND LAURINI, R. *Fundamentals of Spatial Information Systems*. No. 37 in APIC. Academic Press, 1992.
- [31] TRYFONA, N., AND EGENHOFER, M. Consistency among parts and aggregates: A computational model. *Transactions on GIS* 1, 1 (1997), 189–206.

[32] WIJSEN, J. Database repairing using updates. *ACM Trans. Database Syst.* 30, 3 (2005), 722–768.