On First-Order Topological Queries

Martin Grohe, Luc Segoufin

On First-Order Topological Queries

Martin Grohe
Laboratory for Foundations of Computer Science
University of Edinburgh
Edinburgh EH9 3JZ, Scotland, UK
E-mail: grohe@dcs.ed.ac.uk

Luc Segoufin
INRIA-Rocquencourt
B.P. 105,
Le Chesnay Cedex 78153, France
E-mail: Luc.Segoufin@inria.fr

September 25, 2001

Abstract

One important class of spatial database queries is the class of *topological queries*, that is, queries invariant under homeomorphisms.

We study topological queries expressible in the standard query language on spatial databases, first-order logic with various amounts of arithmetic. Our main technical result is a combinatorial characterization of the expressive power of topological first-order logic on *regular* spatial databases.

1 Introduction

The expressive power of first-order logic over finite relational databases is now well understood [Abiteboul et al., Ebbinghaus and Flum, Benedikt and Libkin]. Much less is known in spatial databases (also called constraint databases), where the relations are no longer finite but finitely represented [Kuper et al.].

Most of the expressiveness results in the finite case were obtained by techniques involving *genericity* (invariance under isomorphisms) or *locality* of first-order queries.

The notion of genericity, fundamental for the relational database model, can be generalized to spatial databases in various ways [Paredaens et al.]. Given a group $\mathcal{G}$ of transformations (translations, affinities, isometries, similarities, homeomorphisms, etc.), a query Q is $\mathcal{G}$-generic if for all database instances I and each transformation $g \in \mathcal{G}$, $Q(g(I)) = g(Q(I))$. By $\mathrm{FO}_{\mathcal{G}}$ we denote the set of $\mathcal{G}$-generic first-order queries. The genericity of a first-order query is undecidable [Paredaens et al.] for all interesting groups of transformations, but the expressive power of $\mathrm{FO}_{\mathcal{G}}$ can be understood via *sound and complete (decidable) languages*. A language is said to be *sound* for $\mathcal{G}$ if it contains only $\mathrm{FO}_{\mathcal{G}}$ queries. It is *complete* for $\mathcal{G}$ if it expresses all $\mathrm{FO}_{\mathcal{G}}$ queries. The choice of the group $\mathcal{G}$ depends on which information one is interested in. [Gyssens et al.] gives sound and complete languages for several natural groups of transformations (translations, affinities, isometries, similarities). No sound and complete language is yet known for the group of homeomorphisms.

Queries invariant under homeomorphisms, which are also called *topological queries*, are of fundamental importance in various applications of spatial databases. For example, in geographical databases, queries like "Is region A adjacent to region B?", "Is there a road from A to B?", or "Is A an island?" come up very naturally. Therefore, topological queries have received a lot of attention in the literature (e.g. [Kuijpers et al., Papadimitriou et al., Segoufin and Vianu,

Kuijpers and Van den Bussche]). A basic result known about topological queries is that connectivity of a region is not expressible in first-order logic [Grumbach and Su, Grumbach and Su, Benedikt et al.]. Thinking of geographical databases again, *planar* (or 2-dimensional) database instances, where all relations are embedded in the plane $\mathbb{R}^2$, are of particular importance, and all known results on topological queries concern planar databases.

A first attempt to understand the expressive power of topological queries is to use *topological invariants*. In [Papadimitriou et al.] it has been proven that all topological properties of a planar spatial database can be represented in a finite structure called the *topological invariant* of the instance. In [Segoufin and Vianu] it has been shown how this topological invariant can be used to answer topological queries. In particular, [Segoufin and Vianu] have proven that first-order topological queries on a spatial database can be automatically translated into fixpoint queries on the topological invariant. The translation of first-order topological queries on the spatial database into first-order queries on the topological invariant was proven possible only in the special case of a single relation representing a closed region. It was left open whether this translation could be extended to the case of several regions. We answer this question negatively.

The idea of representing the topological information of a spatial database instance by the topological invariant has two important drawbacks: In a sense, the topological invariant contains too much information; ideally we would just want to store the information that is actually accessible by the query language (which is usually FO). Furthermore, the topological invariant has no straightforward generalization to higher dimensions. The issue of finding an invariant more suitable for FO (and computable in any dimension) was first raised in [Kuijpers et al.].

In the special case of one single relation representing a closed planar region, a *cone structure* was given in [Kuijpers et al.] capturing precisely the first-order topological information. Intuitively, the cone structure is a finite set containing all the possible *small* neighborhoods of a point. The results of [Kuijpers et al.] show that, in this context, first-order topological queries could express only *local* properties, which is a situation known to be true in the finite case. On instances with one closed region that satisfy the additional technical condition of being *fully two dimensional*, [Kuijpers and Van den Bussche] introduced a *cone logic* CL and proved that it is sound and complete for topological FO. [Kuijpers et al.] asked whether their results generalize to database instances with a region that is not necessarily closed or several regions; we give negative answers to these questions.

[Kuijpers et al.] introduced two local operations on planar database instances that preserve the equivalence under first-order topological queries (called *topological elementary equivalence*). We call two instances Ω-*equivalent* if they can be transformed into instances homeomorphic to each other by applying the operations of [Kuijpers et al.] finitely often. Our main technical result, from which all the rest easily follows, is that on especially simple instances that we call *regular*, Ω-equivalence and topological first-order equivalence coincide.

The paper is organized as follows: After recalling a few basic definitions on spatial databases in Section 2, in Section 3 we discuss the topology of planar spatial databases and the topological invariant in detail. In Section 4 we introduce topological first-order queries and review some results of [Kuijpers et al.]. In Section 5, we prove that Ω-equivalence is decidable in PSPACE. Our main result on regular instances is proved in Section 6. In Section 7, we derive that not all first-order topological queries can be translated to first-order queries on the topological invariant, and in Section 8 we briefly discuss the problem of finding a language that is sound and complete for topological FO.

2 Preliminaries

2.1 Spatial databases

We fix an underlying structure $\mathbf{R}$ over the reals; either we let $\mathbf{R} = \mathbf{R}_< = (\mathbb{R}, <)$ or $\mathbf{R} = \mathbf{R}_{\mathrm{lin}} = (\mathbb{R}, <, +, 0, 1)$ or $\mathbf{R} = \mathbf{R}_{\mathrm{poly}} = (\mathbb{R}, <, +, \cdot, 0, 1)$.[1] Let $\sigma_{\mathbf{R}}$ be the vocabulary of $\mathbf{R}$ (i.e. either $\sigma_{\mathbf{R}} = \{<\}$ or $\sigma_{\mathbf{R}} = \{<, +, 0, 1\}$ or $\sigma_{\mathbf{R}} = \{<, +, \cdot, 0, 1\}$).

For a point $\bar{a} = (a_1, \ldots, a_n) \in \mathbb{R}^n$ we let $||\bar{a}|| = \max\{|a_1|, \ldots, |a_n|\}$.[2] For $r \in \mathbb{R}$, let $B_r(\bar{a}) := \{\bar{b} \in \mathbb{R}^n \mid ||\bar{a} - \bar{b}|| < r\}$ be the open ball with radius r around $\bar{a}$, and let $S_r(\bar{a}) := \{\bar{b} \in \mathbb{R}^n \mid ||\bar{a} - \bar{b}|| = r\}$ be the disk with radius r around $\bar{a}$.

A subset $S \subseteq \mathbb{R}^n$, for some $n \geq 1$, is $\mathbf{R}$-*definable*, if there is a first-order formula $\varphi(\bar{x})$ of vocabulary $\sigma_{\mathbf{R}}$ such that $S = \{\bar{a} \in \mathbb{R}^n \mid \mathbf{R} \models \varphi(\bar{a})\}$.

A *schema* is a finite collection σ of *region names*. Let $n \geq 1$. An *n-dimensional spatial database instance* I over σ associates an $\mathbf{R}$-definable set $R^I \subseteq \mathbb{R}^n$ with every $R \in \sigma$. The sets R^I are called the *regions* of I. Formally, we may interpret the region names as n-ary relation symbols and view an instance I over σ as a first-order structure of vocabulary $\sigma_{\mathbf{R}} \cup \sigma$ obtained by expanding the underlying structure $\mathbf{R}$ by the relations R^I, for $I \in \sigma$.

In this paper, we only consider 2-dimensional (or planar) spatial database instances. For convenience, we also assume that all regions are bounded, i.e. that for every instance I over a schema σ and for every $R \in \sigma$ there exists a $b \in \mathbb{R}$ such that $||\bar{a}|| \leq b$ for all $\bar{a} \in R^I$. The boundedness assumption is inessential and can easily be removed, but it occasionally simplifies matters.

2.2 Queries

An *m-ary query* (for an $m \geq 0$) of schema σ is a mapping Q that associates an $\mathbf{R}$-definable subset $Q(I) \subseteq \mathbb{R}^m$ with every instance I over σ. Here we consider $\mathbb{R}^0$ as a one point space and let $\mathbb{R}^0 = \mathrm{TRUE}$, $\emptyset = \mathrm{FALSE}$; a 0-ary query is usually called *Boolean query*. When considering n-dimensional spatial databases, it is reasonable to only admit queries whose arity is a divisible by n.

As our basic query language we take first-order logic FO. A first-order formula $\varphi(x_1, \ldots, x_m)$ of vocabulary $\sigma_{\mathbf{R}} \cup \sigma$ defines the m-ary query

$$I \mapsto \varphi(I) := \{(a_1, \ldots, a_m) \in \mathbb{R}^m \mid I \models \varphi(a_1, \ldots, a_m)\}$$

of schema σ.

3 The topology of planar instances

$\mathbb{R}^2$ is equipped with the usual topology.

The interior of a set $S \subseteq \mathbb{R}^2$ is denoted by $\mathrm{int}(S)$, the closure by $\mathrm{cl}(S)$, and the boundary by $\mathrm{bd}(S)$. We say that a set $S \subseteq \mathbb{R}^2$ *touches* a point $\bar{a} \in \mathbb{R}^2$ if $\bar{a} \in \mathrm{cl}(S)$. Two sets $S_1, S_2 \subseteq \mathbb{R}^2$ *touch* if S_2 touches a point $\bar{a} \in S_1$ or vice versa.

3.1 Stratifications

The existence of stratifications is the fundamental fact that makes the topology of our instances easy to handle.

A *stratification* of an instance I over σ is a finite partition $\mathcal{S}$ of $\mathbb{R}^2$ such that

---

[1] As a matter of fact, we could let $\mathbf{R}$ be any o-minimal structure over the reals, and the main results would remain true.

[2] If $\mathbf{R} = \mathbf{R}_{\mathrm{poly}}$ we may also let $||\bar{a}|| = \sqrt{a_1^2 + \ldots + a_n^2}$, and we will assume we have done so in our figures — it just looks better. Since we are only interested in topological queries, this makes no difference.

(1) For all $S \in \mathcal{S}$, either S is a one point set, or S is homeomorphic to the open interval $(0, 1)$, or S is homeomorphic to the open disk $D = \{\bar{a} \in \mathbb{R}^2 \mid ||\bar{a}|| < 1\}$.

(2) For all $S, S' \in \mathcal{S}$ we either have $\mathrm{cl}(S) \cap \mathrm{cl}(S') = \emptyset$ or $\mathrm{cl}(S) \cap \mathrm{cl}(S')$ is the union of elements of $\mathcal{S}$.

(3) For all $R \in \sigma$ and $S \in \mathcal{S}$ we either have $S \subseteq R^I$ or $S \cap R^I = \emptyset$.

The following lemma follows from the fact that all regions of an instance are **R**-definable. A proof can be found in [van den Dries].

Lemma 3.1 *For every instance I there exists a stratification of I.*

3.2   Colors and cones

Let I be an instance over σ. The *pre-color* of a point $\bar{a} \in \mathbb{R}^2$ is the mapping $\Gamma(\bar{a}) : \sigma \to \{\mathrm{int}, \mathrm{bdi}, \mathrm{bde}, \mathrm{ext}\}$ defined by

$$\Gamma(\bar{a})(R) = \begin{cases} \mathrm{int} & \text{if } \bar{a} \in \mathrm{int}(R^I), \\ \mathrm{bdi} & \text{if } \bar{a} \in \mathrm{bd}(R^I) \cap R^I, \\ \mathrm{bde} & \text{if } \bar{a} \in \mathrm{bd}(R^I) \setminus R^I, \\ \mathrm{ext} & \text{if } \bar{a} \in \mathbb{R}^2 \setminus \mathrm{cl}(R^I). \end{cases}$$

A *pre-cell* is a maximal connected set of points of the same pre-color. The *cone* of a point $\bar{a} \in \mathbb{R}^2$, denoted by $\mathrm{cone}(\bar{a})$, is the (finite) clockwise circular list of the pre-colors of all pre-cells touching $\bar{a}$. Lemma 3.1 implies that cones are well-defined and that there are only finitely many distinct cones in each instance. A point $\bar{a} \in \mathbb{R}^2$ is *regular* if for every neighborhood U of $\bar{a}$ there is a point $\bar{a}' \in U$ such that $\mathrm{cone}(\bar{a}) = \mathrm{cone}(\bar{a}')$. Otherwise $\bar{a}$ is *singular*. It follows from Lemma 3.1 that an instance has only finitely many singular points. We call an instance *regular* if it has no singular points. The cones of regular (singular) points are also called regular (singular, resp.) (cf. Figure


Figure 1: Two singular and two regular cones

1).

The *cone-type* of I, denoted by $\mathrm{ct}(I)$, is a (finite) list of all cones appearing in I. Furthermore, for every singular cone this list also records how often it occurs.

The *color* γ of a point $\bar{a} \in \mathbb{R}^2$ is the pair $\gamma(\bar{a}) = \big(\Gamma(\bar{a}), \mathrm{cone}(\bar{a})\big)$.

3.3   Cells

A *cell* of color γ of I is a maximal connected set of points of color γ. The color of a cell C is denoted by $\gamma(C)$. Lemma 3.1 implies that there are only finitely many cells. Our assumption that all regions are bounded implies that there is precisely one unbounded cell, which we call the *exterior* of I. Lemma 3.1 implies that every cell has a well defined *dimension*, which is either 0, 1, or 2. The 0-dimensional cells are precisely the sets $\{\bar{a}\}$, where $\bar{a}$ is a singular point.

Let $\mathcal{C}^I$ be the set of all cells of an instance I. We define a binary *adjacency* relation $\mathcal{E}^I$ on $\mathcal{C}^I$ by letting two cells be adjacent if, and only if, they touch. We call the graph $\mathcal{G}^I = (\mathcal{C}^I, \mathcal{E}^I)$ the *cell graph* of I. We can partition $\mathcal{C}^I$ into three subsets $\mathcal{C}_0^I$, $\mathcal{C}_1^I$, and $\mathcal{C}_2^I$ consisting of the 0, 1, and 2-dimensional cells, respectively. Observe that the graph $\mathcal{G}^I$ is tri-partite with partition $(\mathcal{C}_0^I, \mathcal{C}_1^I, \mathcal{C}_2^I)$ (i.e. $\mathcal{E}^I \cap (\mathcal{C}_i^I \times \mathcal{C}_i^I) = \emptyset$ for each $i \in \{0, 1, 2\}$). Moreover, $\mathcal{G}^I$ is connected and planar.

4

Lemma 3.2 *Let I be an instance and $C \in \mathcal{C}^I$.*

(1) If $C \in \mathcal{C}_2^I$, then either C is homeomorphic to the open disk $D = B_1((0,0))$, or there exists an $m \geq 1$ such that C is homeomorphic to the open disk D_m with m holes. To be definite, we let

$$D_m := D \setminus \bigcup_{0 \leq i \leq m-1} \operatorname{cl}\left(B_{1/3m}\left(\left(\tfrac{i}{m}, 0\right)\right)\right).$$

(2) If $C \in \mathcal{C}_1^I$, then C is homeomorphic to the sphere $S_1((0,0))$ or to the open interval $(0,1)$.

Proof: This follows easily from Lemma 3.1. □

The *skeleton* $\mathcal{S}^I$ of an instance I is the set of all 0-dimensional cells and all 1-dimensional cells homeomorphic to $(0,1)$. Note that the skeleton of a regular instance is empty.

Lemma 3.3 *Let I be an instance. Then every connected component of the graph $\mathcal{G}^I \setminus \mathcal{S}^I$ is a tree. In particular, if I is regular then $\mathcal{G}^I$ is a tree.*

Proof: Follows from the Jordan Curve Theorem and Lemma 3.2. □

Figure 2 illustrates a typical connected component of an instance I after removing the skeleton. Note that every connected component of the graph $\mathcal{G}^I \setminus \mathcal{S}^I$ has a unique "exterior" cell which we


Figure 2:

may consider as the *root* of the tree. Having the tree directed by fixing this root, we may speak of the *parent* and the *children* of a node.

The following observation, which is a consequence of the fact that the color of a 1-dimensional cell contains the colors of its 2-dimensional neighbors, will be useful later.

Lemma 3.4 *Let I, I' be instances and $C \in \mathcal{C}_1^I, C' \in \mathcal{C}_1^{I'}$ such that $\gamma(C) = \gamma(C')$. Then for every $B \in \mathcal{C}_2^I$ that is adjacent to C there exists a $B' \in \mathcal{C}_2^{I'}$ that is adjacent to C' such that $\gamma(B) = \gamma(B')$.*

3.4 The topological invariant

Two instances I, J over a schema σ are *homeomorphic* if there is a homeomorphism $h : \mathbb{R}^2 \to \mathbb{R}^2$ such that for all $\bar{a} \in \mathbb{R}^2$ and each $R \in \sigma$ we have $\bar{a} \in R^I \iff h(\bar{a}) \in R^J$.

The *topological invariant* of an instance I over σ is an expansion $\mathcal{Y}^I$ of the cell graph that carries enough information to characterize an instance up to homeomorphism. The vocabulary of $\mathcal{Y}^I$ is $\hat{\sigma} := \{E, O, \dim_0, \dim_1, \dim_2, X\} \cup \{\hat{R} \mid R \in \sigma\}$, where E is binary, O is 8-ary, and $\dim_0, \dim_1, \dim_2, X$ and $\hat{R}$ for $R \in \sigma$ are unary. The restriction of $\mathcal{Y}^I$ to $\{E\}$ is the cell graph $\mathcal{G}^I$ of I. $\dim_i$ consists of the i-dimensional cells, for $i = 0, 1, 2$. X only contains the exterior of I (the unique unbounded cell). For every $R \in \sigma$, the unary relation $\hat{R}$ consists of all cells that are subsets of R^I.

O gives the orientation. It is an equivalence relation on the quadruples (C, B, B', B''), where C is 0-dimensional and B, B', B'' are adjacent to C. Two such quadruples

(C_1, B_1, B'_1, B''_1), (C_2, B_2, B'_2, B''_2) are equivalent if either B'_i appears between B_i and B''_i in the clockwise order of the cells adjacent to C_i for both $i \in \{1, 2\}$ or B'_i appears between B_i and B''_i in the anti-clockwise order of the cells adjacent to C_i for both $i \in \{1, 2\}$.[3] Note that O is empty in regular instances.

It is proven in [Papadimitriou et al.] that I and J are homeomorphic if, and only if, $\mathcal{Y}^I$ and $\mathcal{Y}^J$ are isomorphic and that $\mathcal{Y}^I$ is computable from I in time polynomial in the size of I. Since $\mathcal{G}^I$ is a planar graph and canonization of planar graphs is in PTIME (see, e.g., [Grohe]), we can actually assume that $\mathcal{Y}^I$ is *canonical* in the sense that for homeomorphic instances I and J we have $\mathcal{Y}^I = \mathcal{Y}^J$.[4]

4   Topological queries and topological elementary equivalence

4.1   Topological queries

A query is *topological* if for every homeomorphism h of $\mathbb{R}^2$ and for all instances I we have $Q(h(I)) = h(Q(I))$. $\mathrm{FO}_{\mathrm{top}}$ denotes the set of all first-order formulas defining a topological query.

It is well-known (and easy to see) that the set $\mathrm{FO}_{\mathrm{top}}$ is not decidable.

The following lemma collects a few basic FO-queries. Its proof is an easy exercise.

Lemma 4.1   *(1) For every color γ there is a first-order formula $\varphi_\gamma(\bar{x})$ such that for every instance I and for every $\bar{a} \in \mathbb{R}^2$ we have $I \models \varphi_\gamma(\bar{a}) \iff \gamma(\bar{a}) = \gamma$.*

(2) For every $\psi(\bar{y}) \in \mathrm{FO}$ there is a formula $\varphi_{\mathrm{bd}(\psi)}(\bar{x}) \in \mathrm{FO}$ such that for every instance I we have $\varphi_{\mathrm{bd}(\psi)}(I) = \mathrm{bd}(\psi(I))$.

(3) There is a formula $\varphi_\infty(\bar{x}) \in \mathrm{FO}$ such that for every instance I we have

$$\varphi_\infty(I) = \left\{ (a_1, a_2) \in \mathbb{R}^2 \mid \forall R \in \sigma \, \forall (b_1, b_2) \in R^I : |a_1| > |b_1| \text{ and } |a_2| > |b_2| \right\}.$$

Note that for every color γ the formula φ_γ is in $\mathrm{FO}_{\mathrm{top}}$. Moreover, for $\psi \in \mathrm{FO}_{\mathrm{top}}$ the formula $\varphi_{\mathrm{bd}(\psi)}$ is in $\mathrm{FO}_{\mathrm{top}}$. In particular, this is the case for $\psi(\bar{y}) = R(\bar{y})$ for an $R \in \sigma$. On the other hand, the formula φ_∞ is *not* in $\mathrm{FO}_{\mathrm{top}}$.

In [Segoufin and Vianu] it is proven that $\mathrm{FO}_{\mathrm{top}}$-queries over I can be translated in linear time into *fixpoint+counting* queries over the topological invariant $\mathcal{Y}^I$. Furthermore, if $\mathbf{R} = \mathbf{R}_<$ then $\mathrm{FO}_{\mathrm{top}}$-queries over I can be translated to FO-queries over $\mathcal{Y}^I$ on instances with just one closed region. (More precisely, this means that there is a recursive mapping that associates with every $\varphi \in \mathrm{FO}_{\mathrm{top}}$ of vocabulary $\{<, R\}$ a $\varphi' \in \mathrm{FO}$ of vocabulary $\widehat{\{R\}}$ such that for all instances I over $\{R\}$, where R^I is a closed set, we have $I \models \varphi \iff \mathcal{Y}^I \models \varphi'$.)

The question was left open whether this result extends to instances with one arbitrary region or with several regions. In Section 6, we give a negative answer to this question.

4.2   Topological elementary equivalence

Two instances I, J are *(topologically) elementary equivalent* (denoted $I \equiv_t J$) if they satisfy the same topological first-order sentences.

It is proven in [Kuijpers et al.] that if $\sigma = \{R\}$ then for all instances I, J in which R^I, R^J are closed sets we have:

$$I \equiv_t J \iff \mathrm{ct}(I) = \mathrm{ct}(J). \tag{4.1}$$

---

[3] There are various ways to define the orientation, ours is equivalent to [Segoufin and Vianu].

[4] More precisely, there is a PTIME algorithm that, given an instance I, computes $\mathcal{Y}^I$ and a one-to-one mapping $\nu^I : \mathcal{C}^I \to \{1, \dots, |\mathcal{C}^I|\}$ (a *canonical numbering*) such that for homeomorphic instances I, J the mapping $(\nu^J)^{-1} \circ \nu^I : \mathcal{C}^I \to \mathcal{C}^J$ is an isomorphism from $\mathcal{Y}^I$ to $\mathcal{Y}^J$.

We will see in Section 6 that this equivalence cannot be extended to instances with one arbitrary region or with several regions.

To prove this result, [Kuijpers et al.] introduced two simple local operations transforming an instance into an elementary equivalent one. Their straightforward extension to several regions is depicted in Figure 3, which is to be read as follows: Suppose we have an instance I that contains


Figure 3: Operations preserving $\equiv_t$.

an open subset homeomorphic to one of the left hand sides of Figure 3. The different shades of grey display different colors. Then it can be replaced by the corresponding subset on the right hand side (cf. [Kuijpers et al.] for details). Note that both operations are symmetric, we can go from the right to the left by applying the same operation again.

Let ω_1 denote the first and ω_2 the second of the two operations in Figure 3. For instances I and J we write $I \longleftrightarrow_{\omega_i} J$ if I can be transformed into an instance homeomorphic to J by an application of ω_i (for $i \in \{1, 2\}$). We write $I \equiv_\Omega J$ if I and J can be transformed into each other by a finite sequence of operations $\longleftrightarrow_{\omega_1}$ and $\longleftrightarrow_{\omega_2}$.

Then the proof of [Kuijpers et al.] easily yields:

Lemma 4.2 *For all instances I, J we have: $I \equiv_\Omega J \implies I \equiv_t J$.*

It is an open question whether the converse of Lemma 4.2 holds. In particular, this is interesting because is not known whether $\equiv_t$ is decidable or not, whereas $\equiv_\Omega$ is decidable in PSPACE (this is Proposition 5.5 in the next section). [Kuijpers et al.] have shown that $\equiv_t$ and $\equiv_\Omega$ coincide on instances with only one closed region. We can extend their result to several regions, but on instances with only regular cones.

5   Minimal instances

Intuitively a instance is *minimal* if its number of cells cannot be reduced by applying ω_1 or ω_2. More formally we have:

Definition 5.1 *An instance I is* minimal *if it satisfies the following two conditions:*

(M1) *If $C \in \mathcal{C}_1^I$ is homeomorphic to S^1 and $B, B' \in \mathcal{C}_2^I$ are adjacent to C and homeomorphic to D_m, for some $m \geq 1$, then $\gamma(B) \neq \gamma(B')$.*

(M2) *If $B \in \mathcal{C}_2^I$ and $C, C' \in \mathcal{C}_1^I$ are adjacent to B and homeomorphic to $S_1((0,0))$, then $\gamma(C) \neq \gamma(C')$.*

Lemma 5.2 *There is a PTIME algorithm that associates with every instance I a minimal instance $M(I)$ such that $I \equiv_\Omega M(I)$.*

This can be done in such a way that for homeomorphic instances I, J we have $M(I) = M(J)$.

7

Proof: Suppose first that I does not satisfy (M1). We show that I can be transformed to instance J with fewer cells violating (M1), by two applications of ω_1.

Let C be a 1-dimensional cell homeomorphic to S^1 such that both neighbors B_1, B_2 of C have the same color, but neither is homeomorphic to D. Then instance I locally looks like Figure 4(1). We apply ω_1 twice (to the dashed boxes) and obtain an instance that locally looks like Figure 4(3). Note that the 1-dimensional cell in Figure 4(3) does not violate condition (M1), because its interior does not have any holes and thus is not homeomorphic to D_m for an $m \geq 1$. We have obviously reduced the number of cells violating (M1).


(1)   (2)   (3)

Figure 4:

Suppose now that I does not satisfy (M2), then it can be transformed to an instance J with fewer cells violating (M2) by an application of ω_2 without increasing the number of cells violating (M1).

To see this, let B be a 2-dimensional cell in I that is adjacent to the cells C_1, C_2 homeomorphic to S^1 of the same color. A cell homeomorphic to S^1 is adjacent to two 2-dimensional cells. Let B_1, B_2 be the other neighbors of C_1, C_2, respectively. Then by Lemma 3.4, B_1, B_2 also have the same color. We have to distinguish between two cases:

Case 1: Both C_1 and C_2 are children of B. Then we can reduce the number of violations of (M2) by an application of ω_2 (cf. Figure 5).

(1)   (2)

Figure 5:

*Case 2:* C_1 is the parent of B and C_2 its child. Figure 6 shows how to proceed.


(1)   (2)

Figure 6:

A PTIME algorithm transforming a given instance I to a minimal instance $M(I)$ may proceed as follows: Given I, the algorithm first computes the invariant $\mathcal{Y}^I$; this is possible in PTIME

[Papadimitriou et al.]. The operations ω_1 and ω_2 translate to simple local operations on $\mathcal{Y}^I$. Our algorithm first applies pairs of ω_1 until (M1) holds (as in figure 4), and then ω_2 until (M2) holds. This can be done by a simple greedy strategy. The result is a structure $\mathcal{Y}$ that is the topological invariant of a minimal instance $M(I)$. It is shown in [Papadimitriou et al.] that, given an invariant $\mathcal{Y}$, an instance J such that $\mathcal{Y}^J = \mathcal{Y}$ can be computed in polynomial time.

Because $\mathcal{Y}^I$ is canonical (cf. Page 6), this algorithm also guarantees that for homeomorphic instances I, J we have $M(I) = M(J)$. □

5.1 The language of an instance

A *fundamental curve* is an **R**-definable continuous mapping $f : \mathbb{R} \to \mathbb{R}^2$ such that $\lim_{a \to -\infty} ||f(a)|| = \infty$ and $\lim_{a \to \infty} ||f(a)|| = \infty$. We take f to be **R**-definable and not arbitrary because then for every 2-dimensional cell C the set $f^{-1}(C)$ is a finite union of open intervals (one of which may be of the form $(-\infty, a)$ or of the form (b, ∞)), and for every 0- and 1-dimensional cell C the set $f^{-1}(C)$ is a finite union of closed intervals (some of which may just be single points).

We are interested in the finite sequence of colors appearing on a fundamental curve, i.e. in a finite word $W(f, I)$ over the alphabet Σ_I consisting of all colors occurring in I. We say that a word $W \in \Sigma_I^*$ is *realized* in I if there is a fundamental curve f such that $W(f, I) = W$. From the remark above we know that $W(f, I)$ is a finite word. The language $\mathcal{L}(I)$ is the set of all words realized by I.

Example 5.3 *Let $\sigma = \{R\}$ and I be the instance with $R^I := \{\bar{a} \in \mathbb{R}^2 \mid (1/2) < ||\bar{a}|| \leq 1\}$.*

I has five cells $C_1 := \{\bar{a} \in \mathbb{R}^2 \mid ||\bar{a}|| > 1\}$, $C_2 := \{\bar{a} \in \mathbb{R}^2 \mid ||\bar{a}|| = 1\}$, $C_3 := \{\bar{a} \in \mathbb{R}^2 \mid (1/2) < ||\bar{a}|| < 1\}$, $C_4 = \{\bar{a} \in \mathbb{R}^2 \mid ||\bar{a}|| = (1/2)\}$, and $C_5 = \{\bar{a} \in \mathbb{R}^2 \mid ||\bar{a}|| < (1/2)\}$. Let $\alpha, \beta, \gamma, \delta$ be the colors of $C_1, \ldots, C_4$, respectively, and note that C_5 has the same color α as C_1.


Figure 7:

Figure 8:

Then for example the words $\alpha\beta\gamma\delta\alpha\delta\gamma\beta\alpha$ and $\alpha\beta\gamma\beta\alpha\beta\alpha$ are realized in I (cf. Figure 7). It is not hard to see that $\mathcal{L}(I)$ can be described by the finite automaton displayed in Figure 8.

From this example it is easy to see that for every I the language $\mathcal{L}(I)$ is regular; it is accepted by an automaton that is essentially the cell graph.

More formally, let I be an instance. We define a finite automaton $\mathcal{A}^I = (S, \Delta, s_0, F)$ as follows:

- $S := \mathcal{C}^I \cup \{\infty\}$ (where ∞ is a new symbol that does not denote a cell).

- $\Delta := \{(C, \gamma(D), D) \mid C \text{ touches } D\} \cup \{(\infty, \gamma(\text{EXT}), \text{EXT})\}$, where EXT denotes the exterior.

- $F := \{\text{EXT}\}$ and $s_0 := \infty$.

9

Note that the graph underlying $\mathcal{A}^I$ is the cell graph $\mathcal{G}^I$ extended by one additional vertex ∞ that is only adjacent to the exterior. The proof of the following lemma is straightforward:

Lemma 5.4 *For every instance I, the automaton $\mathcal{A}^I$ accepts $\mathcal{L}(I)$. Thus $\mathcal{L}(I)$ is a regular language.*

A *walk* in a graph $G = (V, E)$ is a sequence $\bar{w} := w_1 \dots w_n \in V^n$, for some $n \geq 1$, such that $(w_i, w_{i+1}) \in E$ for $1 \leq i < n$. For a mapping $f : V \to \Sigma$ we let $f(\bar{w}) := f(w_1) \dots f(w_n) \in \Sigma^*$. Then it is almost immediate that for every instance I we have

$$
\mathcal{L}(I) = \{ \quad \gamma(\bar{w}) \mid \bar{w} = w_1 \dots w_n \text{ walk in } \mathcal{G}^I \\
\text{with } w_1 = w_n = \text{EXT} \},
\tag{5.1}
$$

where as usual EXT denotes the exterior.

We are now ready to prove the main result of this section.

Proposition 5.5 $\equiv_\Omega$ *is decidable in* PSPACE.

Proof: Given two instances I and J, we want to check in PSPACE whether it is possible to go from I to J using only homeomorphisms of the plane and operations in Ω. In order to obtain a combinatorial equivalent to homeomorphic transformations we are going to work over the topological invariant of the instances I and J. The operations in Ω have a straightforward equivalent over the structure of the topological invariant that we still call Ω. As proven in [Papadimitriou et al.], I and J are homeomorphic if, and only if, $\mathcal{Y}^I$ and $\mathcal{Y}^J$ are isomorphic. Because topological invariants are essentially planar graphs it is possible to decide in PTIME whether two topological invariants are isomorphic. Given an instance I it is possible to compute its invariant $\mathcal{Y}^I$ in PTIME [Papadimitriou et al.].

From the paragraph above we can conclude that the problem of testing $\equiv_\Omega$ equivalence for instances I and J in PSPACE reduces to checking in PSPACE whether $\mathcal{Y}^J$ can be derived from $\mathcal{Y}^I$ using operations from Ω.

There are at most polynomially many different ways to apply one operator from Ω to $\mathcal{Y}^I$ (one just has to consider tuples of at most 5 cells and check if an ω_i can be applied to this tuple). Let $\Omega(I)$ be the set of invariants that can be obtained from $\mathcal{Y}^I$ by applying one operation from Ω. The previous remarks show that $\Omega(I)$ contains at most polynomially many elements and can be computed in PTIME. It is therefore possible to enumerate all topological invariants that can be derived from $\mathcal{Y}^I$ by applying operations in Ω, and check at each step whether it is isomorphic to $\mathcal{Y}^J$ or not. The latter can be done in PTIME, because it reduces to deciding whether two planar graphs are isomorphic [Hopcroft and Tarjan]. In order to get the PSPACE complexity we show that it is sufficient to consider topological invariants whose size is bounded polynomially in the size of $\mathcal{Y}^I$ and $\mathcal{Y}^J$.

It is easy to see that operations in Ω do not change the cones of the instances involved. Therefore $I \equiv_\Omega J$ implies that I and J have the same cone-type (cf. Page 4) and the same orientation relation O. This can be checked in PTIME. As the cones determine the number of 1-dimensional cells homeomorphic to $(0,1)$, this implies that all instances J such that $I \equiv_\Omega J$ are such that their respective skeletons satisfy $|\mathcal{S}^I| = |\mathcal{S}^J|$.

We may view the skeletons $\mathcal{S}^I$ and $\mathcal{S}^J$ as embedded planar graphs with the 0-dimensional cells as vertices and the 1-dimensional cells homeomorphic to $(0,1)$ as edges. Let F^I and F^J, respectively, denote the faces of these embedded graphs. Note that these faces correspond precisely to the connected components of $\mathcal{G}^I \setminus \mathcal{S}^I$ and $\mathcal{G}^J \setminus \mathcal{S}^J$, respectively. The number of connected components of $\mathcal{S}^J$ is bounded by $c^I < |I|$, the number of cones of I. Therefore the number of faces of $\mathcal{S}^J$ is given by the Euler formula, we have $|F^J| = e - v + 1 + \alpha$, where α is the number of connected components of $\mathcal{S}^J$, and e, v its number of edges and vertices. As α is bounded by $|I|$, so is $|F^J|$. This gives a linear (in $|I|$) bound on the number of connected components of $\mathcal{G}^J \setminus \mathcal{S}^J$.

It would be sufficient to derive the PSPACE complexity to also bound the size of each connected component of $\mathcal{G}^J \setminus \mathcal{S}^J$ (or equivalently the number of 1-dimensional cells homeomorphic to S^1).

Unfortunately, repeated application of the operations ω_1 and ω_2 may produce arbitrary large such components. Nevertheless we are going to show that it is sufficient to consider those that have a size polynomially bounded.

Let K and K' be topological invariants and i be an isomorphism such that $\mathcal{S}^{K'} = i(\mathcal{S}^K)$. We say that $K \equiv_\Omega^i K'$ if K' can be obtained from K using operation of Ω such that all intermediate structure K'' satisfy $\mathcal{S}^{K''} = i(\mathcal{S}^K)$.

We are going to proceed as follows: First we will generate in PSPACE all topological invariants K and isomorphisms i such that $\mathcal{Y}^I \equiv_\Omega K$ and $\mathcal{S}^J = i(\mathcal{S}^K)$. Then for each such K and i we will check in PTIME whether $K \equiv_\Omega^i \mathcal{Y}^J$.

Let us first show that $\equiv_\Omega^i$ can be checked in PTIME.

Claim. Equivalence for $\equiv_\Omega^i$ can be decided in PTIME.

PROOF. Let K and K' be topological invariants and i be a isomorphism such that $\mathcal{S}^{K'} = i(\mathcal{S}^K)$.

By Lemma 5.2 K and K' can be assumed to be minimal.

Recall that every face f of the skeleton F^K of K ($f \in F^{K'}$), corresponds to a connected component $\mathcal{T}_f$ of $\mathcal{G}^K \setminus \mathcal{S}^K$ ($\mathcal{G}^{K'} \setminus \mathcal{S}^{K'}$, respectively). By Lemma 3.3, $\mathcal{T}_f$ is a tree with a canonical root E_f, its exterior cell.

To find out whether $K \equiv_\Omega^i K'$ we check, for each face f of F^K, whether $\mathcal{T}_f$ can be transformed into $\mathcal{T}_{i(f)}$ using operations in Ω that do not affect the skeleton $\mathcal{S}^K$ and the other faces.

Before going on we need the following definitions. A *branch* of a tree $\mathcal{T}$ is a minimal walk in $\mathcal{T}$ going from the root of $\mathcal{T}$ to one of its leaves. A walk in $\mathcal{G}^I$ is said to be *regular* if it never goes through a singular cell. Let Σ be the set of colors of I and J. For a word $W \in \Sigma^*$, we define $M(W)$ as the word of Σ^* computed using rules in the same spirit as in Lemma 5.2. More precisely this means rewriting the word W using the rules $\alpha\beta\alpha \to \alpha$ for all colors α, β.

Fact 1. *Let f be a face in F^K and $\bar{x}$ be a regular walk in $\mathcal{G}^K$ which starts in E_f. Then it is possible to transform K into $K'' \equiv_\Omega^i K$ in such a way that $\mathcal{T}_f$ contains a branch $t_{\bar{x}}$, such that $\gamma(t_{\bar{x}}) = \gamma(\bar{x})$ and such that all other faces of F^K remain unchanged.*

This can be proved by a single loop on the length of $\bar{x} = x_1 \cdots x_n$ starting from the end. Without loss of generality we can assume that x_n is a 2-dimensional cell. We start by constructing in x_{n-2} a new ball $v_{n-1}v_n$ of color $\gamma(x_{n-1}x_n)$. This is easily done by applying ω_1 or ω_2 once. Assume next that we have constructed a subinstance $v_{i+1} \cdots v_n$ whose cell graph is a path attached to x_i. Again we apply ω_1 or ω_2 once in order to get $v_{i-1}v_i$ in x_{i-2} surrounding $v_{i+1} \cdots v_n$.

The same kind of induction (but starting from the beginning of the word this time) shows that the converse also holds:

Fact 2. *Let f be a face of F^K. If it is possible to construct a $K'' \equiv_\Omega^i K$ which adds a new branch t in $\mathcal{T}_f$ then there exists a walk $\bar{x}_t$ in $\mathcal{G}^I$ starting from E_f such that $M(\gamma(\bar{x}_t)) = M(\gamma(t))$.*

We say that a walk $\bar{x}$ *realizes* a word W in I if $M(\gamma(\bar{x})) = M(W)$. Now we can prove the claim. We first modify K as follows. For each face f of F^K we do the following. For each branch t of $\mathcal{T}_{i(f)}$ starting from $E_{i(f)}$, let $W_t = \gamma(t)$ and check whether there is a walk $\bar{x}'$ that realizes W_t (this can be done in PTIME). If this is not the case, Fact 2 shows that $K' \not\equiv_\Omega^i K$. If it is the case, use Fact 1 to construct the corresponding branch in K.

Next we do the same transformation starting from K' and reversing the role of K and K'.

It is clear that if the algorithm does not find out that $K \not\equiv_\Omega^i K'$ on its way, after minimizing the resulting instances we have $\mathcal{T}_f = \mathcal{T}_{i(f)}$ for all $f \in F^K$. As operations in Ω are reversible we can conclude that $K \equiv_\Omega^i K'$.

This proves the claim as all steps above are done in time polynomial in the size of K and K'.

To conclude the proof we will show that in PSPACE we can derive all reachable (by Ω transformations) isomorphism types of the skeleton.

We distinguish between two types of Ω transformations. Type (i) are those that change the skeleton. Type (ii) are those that modify the tree $\mathcal{T}_f$ for some face f of the skeleton.

We are now interested in type (i) transformations. They do not change the size of the topological invariant much: if only type (i) transformations are done, the size of the topological invariant cannot increase or decrease more than the possible number of connected components of the skeleton, which in turn is bounded by the number of cones, which is fixed once for all as mentioned in the beginning of this proof.

Type (ii) transformations can change arbitrarily the size of the topological invariant. Notice that in order to get all possible isomorphism types for the skeleton we need to ensure that connected components of the skeleton interact as much as possible. Two connected components of the skeleton can interact only if they are children of the same node of a tree $\mathcal{T}_f$ for some face f of the skeleton. Increasing the size of a tree $\mathcal{T}_f$ can thus only limit the possible interactions and restrict the investigation of isomorphism types. They can thus be eliminated.

To finish the proof we need the following fact. Let f be a face of the skeleton such that $\mathcal{T}_f$ has two nodes corresponding to connected components S_1 and S_2 of the skeleton.

Fact 3. It can be checked in PTIME whether $\mathcal{T}_f$ can be transformed in such a way that S_1 and S_2 can interact (that is to say they appear as children of the same node).

Assuming $\mathcal{T}_f$ is minimized (if not this can be done in PTIME by Lemma 5.2), the fact is proven by an easy induction on the path from S_1 to S_2 in $\mathcal{T}_f$.

It is now clear how to proceed in the first step. After minimization, Fact 3 is used in order to make connected components of the skeleton interact as much as possible and then all possible isomorphism type are computed in PSPACE.

This concludes the proof of the proposition. $\qquad\square$

In the next section, it will be convenient to work with a slight simplification of the cell graph. We call a 2-dimensional cell $B \neq \text{EXT}$ of an instance I *inessential* if B is homeomorphic to a disk D (and thus has precisely one neighbor in $\mathcal{G}^I$), and the neighbor C of B in $\mathcal{G}^I$ has another neighbor $B' \neq B$ with $\gamma(B') = \gamma(B)$. Let $\mathcal{H}^I$ be the graph obtained from $\mathcal{G}^I$ by deleting all vertices that are inessential cells. We call $\mathcal{H}^I$ the *reduced cell graph* of I. Then (5.1) actually holds with $\mathcal{H}^I$ instead of $\mathcal{G}^I$:

$$\mathcal{L}(I) = \{ \quad \gamma(\bar{w}) \mid \bar{w} = w_1 \ldots w_n \text{ walk in } \mathcal{H}^I$$
$$\text{with } w_1 = w_n = \text{EXT}\}. \tag{5.2}$$

6 Regular instances

Recall that an instance I is *regular* if all points $\bar{a} \in \mathbb{R}^2$ are regular. The main result of this section is that $\equiv_\Omega$ and $\equiv_t$ coincide on regular instances. As a corollary, we will see that the equivalence (4.1) does not extend beyond instances with one closed region.

To illustrate where the problems are, let us start with a simple example:

Example 6.1 *Let* $\sigma = \{R, S\}$ *and consider the two instances* I, J *with* $R^I := \{\bar{a} \in \mathbb{R}^2 \mid \frac{1}{3} \leq \|\bar{a}\| \leq \frac{2}{3}\}$, $S^I := \{\bar{a} \in \mathbb{R}^2 \mid \frac{2}{3} \leq \|\bar{a}\| \leq 1\}$, $R^J := S^I$, *and* $S^J := R^I$ *(cf. Figure 9). Obviously,* I *and* J *have the same cone-type.*


Figure 9:

Let us try to find a sentence $\varphi \in \text{FO}_{\text{top}}$ *such that* $I \models \varphi$ *and* $J \not\models \varphi$. *At first glance this looks easy, just take the sentence saying "every horizontal line that intersects region* R *intersects region*

S before". Then every instance homeomorphic to I satisfies φ, and every instance homeomorphic to J does not satisfy φ.

Unfortunately, $\varphi \notin \mathrm{FO_{top}}$. Figure 10 shows why: All three instances displayed are homeomorphic, but only the last one satisfies φ.

Figure 10:

We will see later that there is a sentence $\varphi \in \mathrm{FO_{top}}$ that distinguishes I from J, but such a φ is quite complicated. For now, let us just note that $I \not\equiv_\Omega J$.

Recall the definition of the minimal instance $M(I)$ associated with an instance I. An inspection the proof of Lemma 5.2 shows that for a regular instance I we have

$$\mathcal{L}(I) \subseteq \mathcal{L}(M(I)). \tag{6.1}$$

Recall that by Lemma 3.3, the cell graph $\mathcal{G}^I$ of a regular instance I is a tree. We think of this tree as being directed with the exterior as its root. The leaves are the 2-dimensional cells homeomorphic to the disk D. Recall the definition of the reduced cell graph $\mathcal{H}^I$ of I. It is also a tree, which is obtained from $\mathcal{G}^I$ by removing those leaves having the same color as their grandparent. Immediately from Definition 5.1 and the fact that cell graphs of regular instances are trees, for minimal regular instances M we get the following:

$$\begin{array}{l} \text{If } C \text{ is a vertex of } \mathcal{H}^M, \text{ then all neighbors of} \\ C \text{ in } \mathcal{H}^M \text{ have different colors.} \end{array} \tag{6.2}$$

For instances I, J we write $\mathcal{H}^I \subseteq_\gamma \mathcal{H}^J$ if there is an (one-to-one) embedding h of $\mathcal{H}^I$ into $\mathcal{H}^J$ that preserves γ. We define $\cong_\gamma$, $\subset_\gamma$ accordingly.

Lemma 6.2 *For all regular instances I, J we have $\mathcal{H}^I \cong_\gamma \mathcal{H}^J$ if, and only if, I and J are homeomorphic.*

Proof: The backward direction is trivial. For the forward direction, let I and J be regular instances with $\mathcal{H}^I \cong_\gamma \mathcal{H}^J$. Then $G^I \cong_\gamma \mathcal{G}^J$. Since for regular instances the orientation O is empty, this implies $\mathcal{Y}^I \cong \mathcal{Y}^J$. Thus I and J are homeomorphic by [Papadimitriou et al.]. $\square$

Lemma 6.3 *Let M, N be minimal regular instances. Then*

$$\mathcal{L}(M) \subseteq \mathcal{L}(N) \iff \mathcal{H}^M \subseteq_\gamma \mathcal{H}^N$$

Proof: This is an easy consequence of (5.2) and (6.2). Indeed take a walk $\bar{w} = w_1 \cdots w_n$ in $\mathcal{H}^M$ going through all its nodes. By (5.2) this walk defines a word in $\mathcal{L}(M)$ which is by hypothesis also in $\mathcal{L}(N)$. By (5.2) this implies that there is a walk $\bar{w}' = w'_1 \cdots w'_n$ in $\mathcal{H}^N$ corresponding to the same word. By (6.2) the function $h : w_i \to w'_i$ is an embedding of $\mathcal{H}^M$ into $\mathcal{H}^N$ preserving γ. $\square$

Recall that a regular language $\mathcal{L} \subseteq \Sigma^*$ is *aperiodic* if there exists an $n \in \mathbb{N}$ such that for all $u, v, w \in \Sigma^*$ such that $uv^n w \in \mathcal{L}$ we also have $uv^{n+1} w \in \mathcal{L}$. By a well-known theorem of Mc-Naughton, Papert [McNaughton and Papert] and Schützenberger [Schützenberger], the aperiodic languages are precisely the languages that are definable in first-order logic.[5]

---

[5]Furthermore, these are precisely the *star-free* regular languages.

The following lemmas will be useful later. They are all based on (6.2) and the fact that the reduced cell graphs of regular instances are trees.

This first lemma states that the language of a minimal instance is aperiodic. This is crucial because we want to capture it by first-order means.

Lemma 6.4 *Let M be a minimal regular instance. Then $\mathcal{L}(M)$ is aperiodic.*

Proof:

Let M be a minimal regular instance.

The crucial step is to prove the following claim:

$$\text{Let } l, n \geq 1, v \in \Sigma^l, \text{ and } \bar{y} = y_1 \ldots y_{nl+1} \text{ a walk in } \mathcal{H}^M \text{ such}$$
$$\text{that } y_1 = y_{nl+1} \text{ and } \gamma(y_{il+1} \ldots y_{(i+1)l}) = v \text{ for } 0 \leq i \leq n-1. \tag{6.3}$$
$$\text{Then } y_1 = y_{l+1}.$$

Suppose for contradiction that (6.3) is wrong. Choose l minimal such that there is a $v = v_1 \ldots v_l \in \Sigma^l$, an $n \geq 2$, and a walk $y_1 \ldots y_{nl+1}$ in $\mathcal{H}^M$ with $y_1 = y_{nl+1}$ and $\gamma(y_{il+1} \ldots y_{(i+1)l}) = v$ for $0 \leq i \leq n-1$, but $y_1 \neq y_{l+1}$.

Since adjacent vertices in a cell graph (and thus in a reduced cell graph) have different colors and $\gamma(y_1) = \gamma(y_{l+1})$, we have $l \geq 2$.

For notational convenience, we let $y_0 = y_{nl}$, $y_{nl+1} = y_1$, $v_0 = v_l$, and $v_1 = v_{l+1}$. We choose an $i \in \{1, \ldots, nl\}$ such that there is no $j \in \{1, \ldots, nl\}$ such that y_j is in the subtree below y_i. Then $y_{i-1} = y_{i+1}$ is the parent of y_i.

Let $\alpha := \gamma(y_i)$, $\beta := \gamma(y_{i-1})$, and $j = i \pmod{l}$. Then $v_{j-1}v_j v_{j+1} = \beta\alpha\beta$. Let $v' := v_1 \ldots v_{j-1}v_{j+2} \ldots v_l$. Then for $0 \leq k \leq n$ we have $\gamma(y_{kl+j-1}) = \beta, \gamma(y_{kl+j}) = \alpha, \gamma(y_{kl+j+1}) = \beta$. By (6.2), this implies $y_{kl+j-1} = y_{kl+j+1}$. If $l = 2$, this implies $y_1 = y_{l+1}$, a contradiction. If $l \geq 3$ we can define a walk $\bar{y}'$ from $\bar{y}$ by deleting y_{kl+j} and y_{kl+j+1}, for $0 \leq k \leq n-1$. Then (6.3) holds with n, $l-2$, v', and $\bar{y}'$, in contradiction to the minimality of l.

This proves (6.3).

Now let $n > |\mathcal{H}^M|$ and $u, v, w \in \Sigma^*$ such that $uv^n w \in \mathcal{L}(M)$. We shall prove that $uv^{n+1}w \in \mathcal{L}(M)$. Let k, l, m be the length of u, v, w, respectively, and $\bar{x} := x_1 \ldots x_{k+nl+m}$ a walk in $\mathcal{H}^M$ with $\gamma(\bar{x}) = uv^n w$.

Since $n > |\mathcal{H}^M|$ there exist i, j with $1 \leq i < j \leq n$ such that $x_{k+il+1} = x_{k+jl+1}$. Applying (6.3) to the walk $\bar{y} := x_{k+il+1} \ldots x_{k+jl+1}$ we infer $x_{k+il+1} = x_{k+(i+1)l+1}$. But then

$$\bar{x}' := \quad x_1 \ldots x_{k+(i+1)l} x_{k+il+1} \ldots x_{k+(i+1)l}$$
$$x_{k+(i+1)l+1} \ldots x_{k+nl+m}$$

is a walk with $\gamma(\bar{x}') = uv^{n+1}w$.

$\square$

Lemma 6.5 *Let M be a minimal regular instance and J, J' instances such that $J \equiv_\Omega J'$ and $\mathcal{L}(J) \subseteq \mathcal{L}(M)$. Then $\mathcal{L}(J') \subseteq \mathcal{L}(M)$.*

Proof: Recall the definitions of the operations ω_1 and ω_2 from Figure 3. It suffices to prove the statement for instances J, J' with $J \longleftrightarrow_{\omega_1} J'$ or $J \longleftrightarrow_{\omega_2} J'$.

To show that $\mathcal{L}(J') \subseteq \mathcal{L}(M)$, let $W' \in \mathcal{L}(J')$. Then there is a word $W \in \mathcal{L}(J)$ obtained from W' by replacing some letters α by subwords $\alpha\beta\alpha$ (if $J \longleftrightarrow_{\omega_1} J'$) or $\alpha\beta\gamma\beta\alpha$ (if $J \longleftrightarrow_{\omega_1} J'$).

Since $\mathcal{L}(J) \subseteq \mathcal{L}(M)$, we have $W \in \mathcal{L}(M)$. Thus there is a walk $\bar{x}$ in $\mathcal{H}^M$ such that $\gamma(\bar{x}) = W$. Let $W = \alpha_1 \ldots \alpha_n$ and $\bar{x} = x_1 \ldots x_n$. By (6.2), whenever $\alpha_i \alpha_{i+1} \alpha_{i+2} = \alpha\beta\alpha$, we have $x_i = x_{i+2}$. Hence $\bar{x}' = x_1 \ldots x_{i-1} x_i x_{i+3} \ldots x_n$ is a walk in $\mathcal{H}^M$ with $\gamma(\bar{x}') = \alpha_1 \ldots \alpha_{i-1}\alpha\alpha_{i+3} \ldots \alpha_n \in \mathcal{L}(M)$. Similarly, if $\alpha_i \ldots \alpha_{i+4} = \alpha\beta\gamma\beta\alpha$, we have $x_i = x_{i+4}$, which implies that $\alpha_1 \ldots \alpha_{i-1}\alpha\alpha_{i+5} \ldots \alpha_n \in \mathcal{L}(M)$. Thus $W' \in \mathcal{L}(M)$. $\square$

This next lemma shows that it is enough to consider horizontal curves. This is important because horizontal curves are definable in first-order.

Lemma 6.6 *Let M be a minimal regular instance and J a regular instance with $\mathcal{L}(J) \nsubseteq \mathcal{L}(M)$. Then there is a $b \in \mathbb{R}$ such that for the curve $f_b : x \mapsto (x, b)$ we have $W(f_b, J) \notin \mathcal{L}(M)$.*

Proof: We assume that J and M are instances over the same schema. Then, denoting by EXT and EXT$'$ the exterior of J, M, respectively, we have $\gamma(\text{EXT}) = \gamma(\text{EXT}')$.

For every walk $\bar{x} = x_1 \ldots x_n$ in $\mathcal{H}^J$ with $x_1 = \text{EXT}$ we define a sequence $\bar{x}' = x_1' \ldots x_n'$ as follows: We let $x_1' := \text{EXT}'$. For $2 \leq i \leq n$ we let x_i' be the (by (6.2) unique) neighbor of x_{i-1}' in $\mathcal{H}^M$ that has the same color as x_i, if $x_{i-1}' \neq \bot$ and such a neighbor exists, and $x_i' := \bot$ otherwise.

Since $\mathcal{H}^J$ is a tree and $\mathcal{H}^M$ satisfies (6.2) for $1 \leq i < j \leq m$ we have

$$x_i = x_j \implies x_i' = x_j' \text{ or } x_j' = \bot. \tag{6.4}$$

It follows that for any walk $\bar{y} = y_1 \ldots y_l$ in $\mathcal{H}^J$ with $y_1 = \text{EXT}$ and $1 \leq i \leq n, 1 \leq j \leq l$ with $x_i = y_j$ we either have $x_i' = \bot$ or $y_j' = \bot$ or $x_i' = y_j'$. To see this, just consider the walk $\bar{x}' = x_1 \ldots x_i x_{i-1} \ldots x_1 y_2 \ldots y_j$ and apply (6.4).

Note that actually the sequence $\bar{x}'$ only depends on the word $\gamma(\bar{x})$. This if $x_n = \text{EXT}$ (and thus $\gamma(\bar{x}) \in \mathcal{L}(J)$), then $\gamma(\bar{x}') \in \mathcal{L}(M)$ if, and only if, $x_n' \neq \bot$.

Now suppose that $\gamma(\bar{x}) \in \mathcal{L}(J) \setminus \mathcal{L}(M)$. Let m be maximal with $x_m' \neq \bot$. Then $m < n$. Observe that x_{m+1} must be a child of x_m in the tree $\mathcal{H}^J$, because if x_{m+1} was the parent of x_m, it would have appeared on the path before entering x_m for the first time, say as x_i, with $x_{i+1} = x_m$. Then by (6.4), $x_{i+1}' = x_m'$, and we could let x_{m+1}' be the parent of x_m'.

Choose $b \in \mathbb{R}$ such that the curve f_b intersects the cell x_{m+1}. Let $\bar{y} = y_1 \ldots y_l$ be a walk in $\mathcal{H}^J$ that corresponds to the curve f_b and let $j < l$ be minimal such that $y_{j+1} = x_{m+1}$. Then $y_j = x_m$, because x_{m+1} is a child of x_m. Then either $y_j' = \bot$, or $y_j' = x_m'$. In both case we have $y_{j+1}' = \bot$. Thus $W(f_b, J) = \gamma(\bar{y}) \notin \mathcal{L}(M)$. $\qquad\square$

Theorem 6.7 *Let I be a regular instance of schema σ. Then there is a sentence $\varphi_I \in \text{FO}_{\text{top}}$ of vocabulary $\{<\} \cup \sigma$ such that an instance J of the same schema σ satisfies φ_I if, and only if, $J \equiv_\Omega I$.*

Proof: Let Σ be the set of colors that may appear in instances of schema σ. For every aperiodic language $\mathcal{L} \subseteq \Sigma^*$ there is a formula $\chi_\mathcal{L}(y) \in \text{FO}$ of vocabulary $\{<\} \cup \sigma$ such that a regular instance I satisfies $\chi_\mathcal{L}(b)$ for a $b \in \mathbb{R}$ if, any only if, $W(f_b, I) \in \mathcal{L}$. This is an easy consequence of the theorem of McNaughton, Papert and Schützenberger that the aperiodic languages are precisely the first-order definable languages.

Let $M := M(I)$ and $\mathcal{L} := \mathcal{L}(M)$. By Lemma 6.4, $\mathcal{L}$ is aperiodic. Let

$$\psi_M := \forall y \chi_\mathcal{L}(y).$$

Clearly, M satisfies ψ_M.

We claim that for all instances J we have:

$$J \models \psi_M \iff M(J) \models \psi_M. \tag{6.5}$$

To prove this claim, note first that every instance satisfying ψ_M realizes the same cones as I and thus is regular.

Let J be a regular instance. Assume first that $J \models \psi_M$. Then by Lemma 6.6, $\mathcal{L}(J) \subseteq \mathcal{L}(M)$. Because $J \equiv_\Omega M(J)$, by Lemma 6.5 we have $\mathcal{L}(M(J)) \subseteq \mathcal{L}(M)$. Thus $M(J) \models \psi_M$.

Conversely, if $M(J) \models \psi_M$, then by Lemma 6.6, $\mathcal{L}(M(J)) \subseteq \mathcal{L}(M)$. By (6.1), it follows $\mathcal{L}(J) \subseteq \mathcal{L}(M)$ and thus $J \models \psi_M$.

This proves (6.5).

It follows easily that $\psi_M \in \text{FO}_{\text{top}}$. Indeed, assume that J satisfies ψ_M. For every J' homeomorphic to J we have $M(J') = M(J)$. Thus $J \models \psi_M$ iff $M(J) \models \psi_M$ iff $J' \models \psi_M$.

We let

$$\varphi_I := \psi_M \wedge \bigwedge_{\substack{N \text{ minimal regular instance} \\ \text{with } \mathcal{H}^N \subset_\gamma \mathcal{H}^M}} \neg\psi_N.$$

Then $\varphi_I \in \mathrm{FO}_{\mathrm{top}}$.

We shall prove that for any instance J of schema σ we have

$$J \models \varphi_I \iff J \equiv_\Omega I.$$

Suppose that $J \models \varphi_I$. Then in $M(J) \models \varphi_I$ by (6.5). By Lemma 6.6, this implies $\mathcal{L}(M(J)) \subseteq \mathcal{L}(M)$ and $\mathcal{L}(M(J)) \not\subseteq \mathcal{L}(N)$ for any minimal regular N with $\mathcal{H}^N \subset_\gamma \mathcal{H}^M$. By Lemma 6.3, this implies $\mathcal{H}^{M(J)} \cong_\gamma \mathcal{H}^M$. By Lemma 6.2, it follows that M and $M(J)$ are homeomorphic. Thus $J \equiv_\Omega I$.

Conversely, suppose that $J \equiv_\Omega I$. Then by Lemma 4.2, $J \equiv_t I$. Thus $J \models \varphi_I$, because $\varphi_I \in \mathrm{FO}_{\mathrm{top}}$ and $I \models \varphi_I$ (the later follows easily from the previous paragraph and (6.5)). $\qquad\square$

Corollary 6.8 *For all regular instances I, J we have $I \equiv_\Omega J \iff I \equiv_t J$.*

Finally, we are ready to prove that the equivalence (4.1) does not extend beyond instances with one closed region.

Corollary 6.9 *The two instances of Example 6.1 are not elementary equivalent. Neither are the instances I, J over $\{R\}$ defined by:*

$$R^I := \{\bar{a} \in \mathbb{R}^2 \mid \frac{1}{2} < ||\bar{a}|| \le 1\},$$

$$R^J := \{\bar{a} \in \mathbb{R}^2 \mid \frac{1}{2} \le ||\bar{a}|| < 1\}$$

Proof: Both instances I and J are regular minimal instances. It is easy to see that the trees $\mathcal{G}^I$ and $\mathcal{G}^J$ are not isomorphic and thus by Lemma 6.3 the languages $L(I)$ and $L(J)$ are different and by Lemma 6.5 $I \not\equiv_\Omega J$. Therefore by Corollary 6.8 they are not elementary equivalent. $\qquad\square$

7 Translating sentences to the topological invariant

Recall that it is proven in [Segoufin and Vianu] that there is a recursive mapping that associates with every $\varphi \in \mathrm{FO}$ of vocabulary $\{<, R\}$ a $\varphi' \in \mathrm{FO}$ of vocabulary $\{\hat{R}\}$ such that for all instances I over $\{R\}$, where R^I is a closed set, we have $I \models \varphi \iff \mathcal{Y}^I \models \varphi'$.

The purpose of this section is to prove that this does not extend to arbitrary instances.

Proposition 7.1 *There is a sentence $\varphi \in \mathrm{FO}_{\mathrm{top}}$ of vocabulary $\{<, R\}$ such that for every sentence $\varphi' \in \mathrm{FO}$ of vocabulary $\{\hat{R}\}$ there is an instance I such that $I \models \varphi$ and $\mathcal{Y}^I \not\models \varphi'$.*

Proof: Let I_0, J_0 be the instances defined by

$$R^{I_0} := \{\bar{a} \mid 1 \le ||\bar{a}|| < 2\},$$

$$R^{J_0} := \{\bar{a} \mid 1 < ||\bar{a}|| \le 2\}.$$

For $n \ge 1$, let I_n, J_n be defined by

$$R^{I_n} := R^{I_0} \cup \bigcup_{i=1}^{n} S_{i/n+1}((0,0)) \cup \bigcup_{i=1}^{n} S_{2+i/n+1}((0,0)),$$

$$R^{J_n} := R^{J_0} \cup \bigcup_{i=1}^{n} S_{i/n+1}((0,0)) \cup \bigcup_{i=1}^{n} S_{2+i/n+1}((0,0)),$$

Figure 11: The instances I_3 and J_3.

where $S_r(\bar{a}) := \{\bar{b} \mid ||\bar{b} - \bar{a}|| = r\}$ for all $r \in \mathbb{R}, \bar{a} \in \mathbb{R}^2$ (cf. Figure 11).

Note that for all $n, m \geq 1$ we have $I_n \equiv_\Omega I_m$ and $J_n \equiv_\Omega J_m$, but $I_n \not\equiv_\Omega J_n$ (see the proof of Corollary 6.9). Corollary 6.8 implies that there is a sentence $\varphi \in \mathrm{FO_{top}}$ such that for all $n \geq 1$ we have $I_n \models \varphi$ but it can be checked that $J_n \not\models \varphi$.

Let $n \geq 1$. The graph $\mathcal{G}^{I_n}$ is just a path with $5 + 4n$ vertices, say, $C_1, \ldots, C_{4n+5}$. Denoting the colors by $\alpha, \ldots, \varepsilon$, the colors on this path form the following sequence:

$$\alpha \underbrace{\beta\alpha \ldots \beta\alpha}_{n \text{ times}} \gamma\delta\varepsilon \underbrace{\alpha\beta \ldots \alpha\beta}_{n \text{ times}} \alpha.$$

$\mathcal{G}^{J_n}$ is the same, except that $\gamma(C_{2n+2}) = \varepsilon$ and $\gamma(C_{2n+4}) = \gamma$. The other relations of the topological invariant are identical in $\mathcal{Y}^{I_n}$ and $\mathcal{Y}^{J_n}$. $\dim_0$ is empty, $\dim_1$ consists of all C_i with even i, and $\dim_2$ consists of the C_i with odd i. The orientation O is empty. Finally, $X = \{C_1\}$.

Standard Ehrenfeucht-Fraïssé techniques show that for every sentence $\varphi' \in \mathrm{FO}$ there is an $n \geq 1$ such that $\mathcal{Y}^{I_n} \models \varphi'$ if, and only if, $\mathcal{Y}^{J_n} \models \varphi'$.

The statement of the proposition follows. $\qquad\qquad\square$

8   On completeness of languages

An open problem that we have not considered so far is to find a (recursive) language that expresses precisely the first-order topological queries. Although this is certainly an interesting question, we doubt that, even if theoretically such a language exists, it would be a natural language that may serve as a practical query language. Our results show that first-order topological queries are not local; on the other hand, it is known that first-order logic fails to express the most natural non-local topological queries such as connectivity of a region.

[Kuijpers and Van den Bussche] have introduced a topological query language, the *cone-logic* CL, only expressing local properties. This language is a two tier language that allows to build first-order expressions whose atoms are again first-order expressions talking about the cones and colors of points. [Kuijpers and Van den Bussche] have proven that CL captures precisely the first-order topological properties of instances with only one closed region that is "fully 2-dimensional". If the underlying structure is $\mathbf{R}_<$, it follows from [Segoufin and Vianu] that the condition of being full 2-dimensional is not needed. Corollary 6.9 shows that this result does not extend to instances with several closed regions or one arbitrary region.

We propose to extend CL by a path operator, as it has been introduced in [Benedikt et al.]. Let us call the resulting topological query language PCL. The results of [Benedikt et al.] show that this language has the basic properties expected from a reasonable spatial database query language. In addition, it admits efficient query evaluation (the cost of evaluating a PCL-query is not substantially higher than the cost of evaluating a first-order query).

An example of a PCL-query not expressible in FO is connectivity of regions. We conjecture that conversely every $\mathrm{FO_{top}}$-query is expressible in PCL. The idea behind this conjecture is that local properties (expressible in CL) together with the language of an instance, which can be described by the path-operator, seem to capture the first-order topological properties of an instance. As a first step towards proving this conjecture, let us remark that Corollary 6.8 implies that on regular instances every $\mathrm{FO_{top}}$-sentence is equivalent to a set of PCL-sentences.

9 Conclusions

The results of this paper give a good understanding of first-order topological queries on regular instances. Of course one may argue that regular instances are completely irrelevant — look at any map and you will find singular points. However, we could use our results to answer several open questions concerning topological queries on arbitrary instances.

As a matter of fact, we have shown that the previous understanding of first-order topological queries, viewing them as "local" in the sense that they can only speak about the colors of points, is insufficient; this may be our main contribution.

The problem of a characterization of topological elementary equivalence on arbitrary planar instances remains open. We conjecture that Corollary 6.8 generalizes, i.e. that $\equiv_t$ is the same as $\equiv_\Omega$ on arbitrary instances. If this was true, by Proposition 5.5 topological elementary equivalence would be decidable in PSPACE. Let us remark that we do not believe that the PSPACE-bound of Proposition 5.5 is optimal, we see no reason why $\equiv_\Omega$ should not be decidable in NP or even PTIME.

References

[Abiteboul et al.] ABITEBOUL, S., HULL, R., AND VIANU, V. 1995. *Foundations of Databases*. Addison-Wesley.

[Benedikt et al.] BENEDIKT, M., DONG, G., LIBKIN, L., AND WONG, L. 1998. Relational expressive power of constraint query languages. *Journal of the ACM 45*, 1–34.

[Benedikt et al.] BENEDIKT, M., GROHE, M., LIBKIN, L., AND SEGOUFIN, L. 2000. Reachability and connectivity queries in constraint databases. In *Proceedings of the 19th ACM Symposium on Principles of Database Systems*. 104–115.

[Benedikt and Libkin] BENEDIKT, M. AND LIBKIN, L. 2000. Expressive power: The finite case. In *Constraint Databases*, G. Kuper, L. Libkin, and J. Paredaens, Eds. Springer-Verlag, Chapter 3, 55–88.

[Ebbinghaus and Flum] EBBINGHAUS, H.-D. AND FLUM, J. 1995. *Finite Model Theory*, second ed. Springer-Verlag.

[Grohe] GROHE, M. 1998. Fixed-point logics on planar graphs. In *Proceedings of the 13th Annual IEEE Symposium on Logic in Computer Science*.

[Grumbach and Su] GRUMBACH, S. AND SU, J. 1997a. Finitely representable databases. *Journal of Computer and System Sciences 55*, 273–298.

[Grumbach and Su] GRUMBACH, S. AND SU, J. 1997b. Queries with arithmetical constraints. *Theoretical Computer Science 173*, 151–181.

[Gyssens et al.] GYSSENS, M., VAN DEN BUSSCHE, J., AND VAN GUCHT, D. 1997. Complete geometrical query languages. *Journal of Computer and System Sciences 58*, 483–511.

[Hopcroft and Tarjan] HOPCROFT, J.E. AND TARJAN, R. 1972. Isomorphism of planar graphs. In *Complexity of Computer Computations*. Plenum Press.

[Kuijpers et al.] KUIJPERS, B., PAREDAENS, J., AND VAN DEN BUSSCHE, J. 2000. Topological elementary equivalence of closed semi-algebraic sets in the real plane. In *The Journal of Symbolic Logic*, 65: (4), 1530-1555,

[Kuijpers and Van den Bussche] KUIJPERS, B. AND VAN DEN BUSSCHE, J. 1999. On capturing first-order topological properties of planar spatial databases. In *Proceedings of the 7th International Conference on Database Theory*, C. Beeri and P. Buneman, Eds. Lecture Notes in Computer Science, vol. 1540. Springer-Verlag, 187–198.

[Kuper et al.] Kuper, G., Libkin, L., and Paredaens, J., Eds. 2000. *Constraint Databases*. Springer-Verlag.

[McNaughton and Papert] McNaughton, R. and Papert, S. 1971. *Counter-free Automata*. MIT Press.

[Papadimitriou et al.] Papadimitriou, C., Suciu, D., and Vianu, V. 1999. Topological queries in spatial databases. *Journal of Computer and System Sciences 58*, 29–53.

[Paredaens et al.] Paredaens, J., Van Den Bussche, J., and Van Gucht, D. 1994. Towards a theory of spatial database queries. In *Proceedings of the 13th ACM Symposium on Principles of Database Systems*. 279–288.

[Schützenberger] Schützenberger, M. 1965. On finite monoids having only trivial subgroups. *Information and Control 8*, 190–194.

[Segoufin and Vianu] Segoufin, L. and Vianu, V. 2000. Querying spatial databases via topological invariants. In *Journal of Computer System Sciences 61(2)*, 270–301.

[van den Dries] van den Dries, L. 1998. *Tame Topology and O-minimal Structures*. Cambridge University Press.