The Logic in Computer Science Column

BY

Yuri Gurevich

Microsoft Research
One Microsoft Way, Redmond WA 98052, USA
gurevich@microsoft.com

First-order topological properties

Jan Van den Bussche
Universiteit Hasselt

Author: Oh, hi, you're Quisani, Yuri Gurevich's student, right?

Quisani: That's me allright; can I help you?

A: Well, yes, Yuri asked me if I wanted to write a logic column, and suggested to talk to you for some inspiration.

Q: Sure, we can talk. Do you already have an idea what you will write about?

A: I was thinking, perhaps I could write about first-order topological properties.

Q: I suppose that is first-order as in "first-order logic", but what's about the topology? I'm afraid I do not know more about topology than what you find about it in a dictionary. It's supposed to deal with the shapes of objects; things that do not change when you deform an object, right?

A: That's quite right. Topology is basically that part of geometry that deals with properties that remain invariant under continuous transformations.

Q: But what has that to do with logic?

A: In itself not much. But it is a legitimate question to ask which properties of geometrical objects are at the same time topological and expressible in first-order logic.

Q: Why would anyone care about that?

A: Well, the original motivation comes from the theory of spatial databases.

Q: Do you mean databases the NASA folks use?

A: They probably do use spatial databases at NASA, but the general term does not specifically refer to space in that sense. A spatial database is simply any database that contains data with a geometrical interpretation [24]. Typical application areas are Geographical Information Systems (GIS) and robotics [29, 8].

Q: I see. I suppose topology could indeed be relevant in such applications.

A: You bet. For certain applications, only the topology of the data is important, and not the more specific geometrical aspects like precise locations or rotations. Applications concerned with dimensionality, or with connectivity, fall in this category.

Q: Could you be a little bit more concrete?

A: Sure, take the Paris metro map for example. It gives you full information about the topology of the subway network: how stations are connected. But it is very unreliable in other aspects, such as precise distances.

Q: And people are fine with that, because anyway they use the map only to check how they can get from point A to point B, and are happy with using the number of intermediate stations on the route as a purely topological approximation of distance. OK, I'm with you. But where does the first-order logic come in the picture?

A: We're getting there. Let's first agree that we are interested in the topology of some geometrical object, which we formalize in the standard mathematical way as a set A of points in n-dimensional real space $\mathbb{R}^n$.

A: By using Cartesian coordinates, we can view A as an n-ary relation on $\mathbb{R}$.

Q: Excellent remark, because that's exactly what we will do. This allows us to use first-order logic sentences φ over the vocabulary $(<, 0, 1, +, \times, S)$, evaluated over $\mathbb{R}$, to express properties about geometric objects. Here, $<$, 0, 1, + and $\times$ are the obvious predicates and functions over $\mathbb{R}$, and S is an n-ary relation symbol interpreted by the set A we want to talk about.

Q: OK, let's see… if I want to express that $S \subseteq \mathbb{R}^2$ is a straight line, I can write

$$\exists a \exists b \forall x \forall y (S\,xy \leftrightarrow a \times x + b = y).$$

A: Very good. It is not a topological property, but it is a first-order expressible property allright. To give you an example of a first-order property that is topological, consider the property that $S \subseteq \mathbb{R}^2$ contains a two-dimensional subset. For that we can write

$$\exists x_0 \exists y_0\, \exists r > 0\, \forall x \forall y ((x - x_0)^2 + (y - y_0)^2 < r \rightarrow S\,xy).$$

Q: OK. Now let me try to express that S is topologically connected. For $n = 1$ that is easy, because a subset of the real line is connected if and only if it is an open, half-open, or closed interval, and we can easily express that in first-order logic. But I do not immediately see how to do it in $\mathbb{R}^2$.

A: You will never see it, because it is impossible to do. This was shown by the combined results of Grumbach and Su [15] and Benedikt, Dong, Libkin and Wong [3].

Q: But doesn't that kill your whole story? I thought connectivity questions provided the motivation to study this stuff in the first place.

A: I didn't say that. I said that connectivity is a typical example of a topological question, but there are many others. First-order logic is the theoretical foundation for all database query languages [1], including spatial database query languages, and therefore it is important that we understand precisely which topological properties are expressible in first-order logic, even if we already know connectivity is not one of them.

Q: Fair enough. I see you are eager to tell me what those expressible properties are, but please allow me one further question first. Now that you tell me you are considering first-order logic to be a query language, I wonder how can one evaluate first-order logic sentences over arbitrary, typically infinite subsets of $\mathbb{R}^n$ in an effective manner? Clearly, we want our query language to be implementable on a computer, don't we?.

A: Excellent point; I have almost forgotten to tell you about that. You know about the decidability of the first-order theory of the reals?

Q: Sure, this is an old theorem of Tarski [27]; he showed that the truth of any first-order sentence about the reals can be effectively determined. In the setting we are talking about, these would be the sentences without the extra predicate S. The example they always give of a true sentence over the reals is the solution of quadratic equations:

$$\forall a \forall b \forall c \exists x (ax^2 + bx + c = 0 \leftrightarrow b^2 - 4ac \geq 0)$$

A: Very good. In view of this decidability, we will restrict attention to sets A that are first-order definable over $\mathbb{R}$. That allows us to effectively evaluate a query φ on a set A, simply by plugging in the formula that defines A at all places in φ where the predicate symbol S is used.

Q: But are the sets in $\mathbb{R}^n$ that are explicitly definable by a first-order logic formula over $\mathbb{R}$ interesting enough in practice?

A: They sure are. They include all the geometric objects one encounters in elementary geometry, and form a well-studied class of sets known as the "semi-algebraic sets". They are sufficient for robotics applications [25] and are surely enough for GIS and computer graphics applications, where all the data objects are typically built up from straight line segments.

Q: OK, OK, I see; descriptions of geometrical objects in Cartesian coordinates using polynomial equations or inequalities, such as lines, arcs, circles, cubes, ellipsoids, cylinders, and so on, all fall in the realm of first-order logic formulas. And then we can also take projections (through existential quantification), unions, intersections, complements.

But, isn't going through the first-order theory of the reals an enormously inefficient way of implementing your query language?

A: Of course, specific operations on specific geometric objects can be implemented much more efficiently using specific algorithms from computational geometry. It is the typical trade-off between specificity and generality. Note, however, that algorithmic progress on decision procedures for the reals has been ongoing ever since Collins's cylindrical algebraic decomposition method [2, 7, 16]. And if nothing else, this whole idea of "plug-in evaluation" remains a nice theoretical framework. It was first proposed by Kanellakis, Kuper and Revesz in the form of "constraint databases" [17, 21]. Peter Revesz and his students have implemented quite a few constraint database systems, for different logical theories.

Q: Very interesting. But perhaps we should move on.

A: Yes, let's move on. What I wanted to show you is a characterisation of the first-order topological properties of closed semi-algebraic sets in the plane.

Q: So we're focusing here on subsets of the plane, $\mathbb{R}^2$, but what does "closed" mean?

A: It's a standard topological concept. It means that the set includes all its border. For example, the set of points inside the unit circle, $x^2 + y^2 < 1$, is not closed, but the points inside together with the circle itself, $x^2 + y^2 \leq 1$, is closed. Actually, you can express that S is closed in first-order:

$$\forall x \forall y (\mathrm{Acc}(x, y) \rightarrow S\,xy),$$

where $\mathrm{Acc}(x, y)$ stands for the formula

$$\exists \varepsilon > 0\, \forall \delta \in\,]0, \varepsilon[\; \exists x_0 \exists y_0 (S\,x_0 y_0 \wedge 0 < (x - x_0)^2 + (y - y_0)^2 < \delta)$$

Q: Got it; $\mathrm{Acc}(x, y)$ expresses that (x, y) is an accumulation point. Anyway, from a practical standpoint, restricting to closed sets does not seem too harmful. I wouldn't know how to draw a set on a piece of paper without drawing its borders as well!

A: You could use different colors, e.g., draw the borders that are not part of the set in red and the rest in green. But for a drawing in one color you're right.

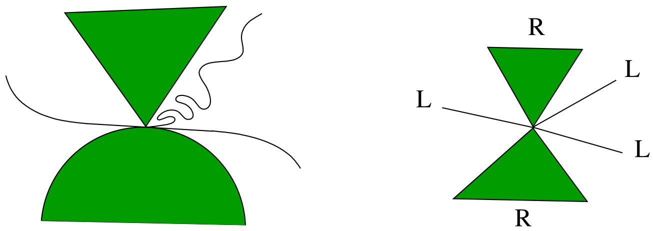

Figure 1: A set and the cone of one of its points.

Now locally around each of its points, a closed semi-algebraic set in the plane looks "conical", in the sense that if you travel around the point in a small enough radius, you will always encounter the same circular list of lines and regions. We call that circular list of L's and R's the "cone" of the point. Let me illustrate it with a drawing (Figure 1).

Q: I see. Is this so because the set is closed?

A: Not really, it is rather because the set is semi-algebraic. Semi-algebraicity rules out "wild" topologies [28]. If the set is not closed, the cones may be a bit more complicated than the ones we have here.

Q: Let me look at some special cases to make sure I understand your definition of cone. If I have a point lying one a line, then its cone is (*LL*), because we see the line to the left and the right of the point. For an endpoint of a line, the cone is simply (*L*). And for a point in the interior of the geometrical object, the cone is (*R*), right?

A: You're right about the points on a line, and about the endpoints, but a point with cone (*R*) is a point on the border of a region. Actually, an interior point is completely surrounded by a region, something we do not have a notation for yet. So let's introduce one and indicate the cone of an interior point by the special letter *F* (for "full").

Q: OK. So in your drawing (Figure 1), all the different cones that we can see, apart from (*RLLRL*) for the central point, are three times (*L*) for the three endpoints of lines; and infinitely many (*LL*)'s, (*R*)'s, and *F*'s, for all the points on the lines, on the borders of the regions, and inside the regions, respectively.

Now that I think about it: the cones (*LL*), (*R*) and *F* always occur infinitely often if they occur at all. Moreover, cones *F* occur if and only if at least one of the cones has an *R*. So, *F* is in a sense redundant.

A: Good observation, unless the set consists of the entire plane, but let's forget about that special case. Now I can present you with a first theorem [20]: two sets in which precisely the same cones appear, with precisely the same multiplicities,

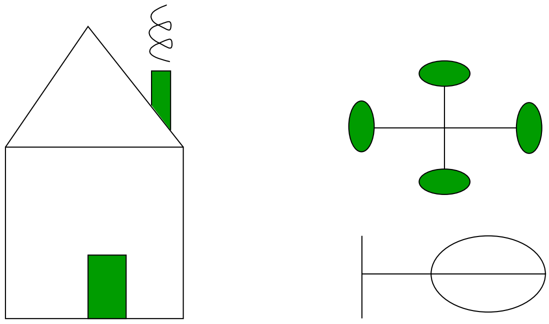

Figure 2: The set on the left has precisely the same cones, with precisely the same multiplicities, as the set on the right. Hence the two sets are indistinguishable by topological first-order sentences.

are indistinguishable by topological first-order sentences. Here's an illustration (Figure 2).

Q: Awesome! How can such a thing be proved?

A: You should look at the paper [20], but in brief, a number of elementary transformations are introduced by which two sets with the same cones can be transformed into each other. These transformations are shown to be indistinguishable by topological first-order sentences. The proof technique involves a reduction to finite structures [15], and an fundamental tool is provided by the collapse theorems for embedded finite models [3, 5, 22].

Q: I'll add the chapter on embedded finite models of Libkin's book to my bedtime reading list.

A: Sweet dreams! Anyway, this indistinguishability theorem really paves the road for a full characterisation of the first-order topological properties.

Q: Hate to stop you now, but one quick remark: I suppose the indistinguishability theorem is actually an if and only if, right? I mean, it seems a doable exercise to express in first-order that a point has a given cone, so if two sets do not agree on their cones then they are distinguishable by a topological first-order sentence.

A: You got it. Now our second theorem lifts the indistinguishability theorem to the global level of properties, and says that the first-order topological properties are precisely the properties that can be expressed in first-order logic when looking only at the cones.

Q: To make sense of that statement we need some kind of logical interpretation of cones.

A: Precisely. But that is very natural. Recall that a cone is a circular list of L's and R's. We can represent such a list as a finite structure $\{1, 2, \ldots, n\}$, where n is the length of the list, equipped with the following relations: L and R are unary relations containing the positions that are L and R, respectively; B is the ternary relation consisting of the triples (i, j, k) of distinct positions such that j comes before k in the sequence

$$i, i + 1, \ldots, n, 1, 2, \ldots, i - 1.$$

Q: So B is a "circular" order relation; I assume 'B' stands for 'between'?

A: You got it. We can now use first-order logic sentences over the vocabulary (L, R, B) to express properties of cones.

Q: Let me try. To express that there are two consecutive L's in the cone, I could write

$$\exists x \exists y \neg \exists z (Bxzy \wedge Lx \wedge Ly \wedge x \neq y)$$

A: Correct! We are now ready to define "cone logic sentences". These are simply boolean combinations of basic sentences of the form $|\gamma| \geq k$, where γ is a first-order sentence over (L, R, B) as above, and k is a natural number. The meaning of such a basic sentence is simply that there are at least k points in the set whose cone satisfies γ.

Q: I see. So, to express that the set is a bunch of non-intersecting lines and loops, i.e., that the only cones that can occur are (L) and (LL), we could write

$$\neg(|\gamma| \geq 1)$$

where γ is

$$\exists x \exists y \exists z (x \neq y \wedge x \neq z \wedge y \neq z) \vee \exists x \, Rx$$

A: The formal theorem now is that the first-order topological properties of closed semi-algebraic sets in the real plane are precisely those that can be expressed as cone logic sentences [6]. Not surprisingly, it is undecidable whether a given first-order sentence is topological. So, the theorem gives us an alternative, syntactic characterisation of an undecidable class of sentences.

Q: I can sympathize with your enthusiasm. It is a beautiful result. How is it proved?

A: You should again look in the paper (a full version is available from me), but in brief, sets are put in a normal form consisting of drawings of cones. Sets in this normal form can be represented by abstract finite structures, which we call "codes". We then show how to rewrite a topological first-order sentence φ into a first-order sentence ψ about codes. The major technical hurdle left after that is that the codes contain information about how cones are linked together by lines. Because we know that this linking information is not captured by φ, we can remove it also from ψ, but this requires a few complicated invariance arguments. The resulting ψ sans linking information then easily yields a cone logic sentence.

Q: Sounds like a nice achievement; congratulations!

A: It was certainly not my achievement alone; a lot of credit goes to my collaborators Michael Benedikt, Bart Kuijpers, Christof Löding, and Thomas Wilke.

Q: So, what does the future hold?

A: A lot of remaining open questions. First of all, we now have a characterisation of the first-order topological properties, but this is for closed semi-algebraic sets in the plane only. What about non-closed sets? A semi-algebraic set can always be written as a boolean combination of closed semi-algebraic sets. Hence, we can ask more generally, what about properties not of a single set, but of a collection of sets, even closed sets? Grohe and Segoufin [14] have shown that already the indistinguishability theorem fails for non-closed sets, even within a class of very simple sets. For that class they do give a new indistinguishability theorem, however.

Further, what in higher dimensions? And what about properties of more general, not necessarily semi-algebraic, sets? On the one hand, the problem becomes more difficult, because we lose the tame topology of semi-algebraic sets. On the other hand, the problem becomes easier, because less first-order sentences will be topological now.

Q: That's quite a research program you have there!

A: I must admit, though, that only a handful of people have worked on this topic. Apart from the people whose papers I've already cited, there are Papadimitriou, Suciu and Vianu [23] who investigated logics over the plane graph representation of the topology of the set, and Segoufin and Vianu [26], who continued that line of research with some very interesting results. As far as I know, currently none of these people is very much occupied with trying to continue the topic. We need new blood!

Q: We're not going to end our conversation on a pessimistic note, are we?

A: No, we can't do that, can we! At any rate, it is quite clear already now that first-order logic is much too weak to be a useful topological query language. So, rather than invest effort in understanding better exactly how weak it is, that effort is perhaps better spent on the investigation of extensions of the language. There has already been interesting work in that direction [4, 13, 10, 18, 19, 11, 12]. Interestingly, a good understanding of extensions of the first-order language can necessitate further work on the first-order language itself [9].

Q: Whew. I have plenty to read and think about now.

A: You've been a great help to me in the writing of this column.

Q: You're welcome.

References

[1] S. Abiteboul, R. Hull, and V. Vianu. *Foundations of Databases*. Addison-Wesley, 1995.

[2] S. Basu, R. Pollack, and M.-F. Roy. *Algorithms in Real Algebraic Geometry*. Springer, 2003.

[3] M. Benedikt, G. Dong, L. Libkin, and L. Wong. Relational expressive power of constraint query languages. *Journal of the ACM*, 45(1):1–34, 1998.

[4] M. Benedikt, M. Grohe, L. Libkin, and L. Segoufin. Reachability and connectivity queries in constraint databases. *Journal of Computer and System Sciences*, 66(1):169–206, 2003.

[5] M. Benedikt and L. Libkin. Relational queries over interpreted structures. *Journal of the ACM*, 47(4):644–680, 2000.

[6] M. Benedikt, C. Löding, J. Van den Bussche, and T. Wilke. A characterization of first-order topological properties of planar spatial data (extended abstract). In *Proceedings 23th ACM Symposium on Principles of Database Systems*, pages 107–114. ACM Press, 2004.

[7] B.F. Caviness and J.R. Johnson, editors. *Quantifier Elimination and Cylindrical Algebraic Decomposition*. Springer, 1998.

[8] M. de Berg, M. van Kreveld, M. Overmars, and O. Schwarzkopf. *Computational Geometry*. Springer, 2000.

[9] F. Geerts. Expressing the box cone radius in the relational calculus with real polynomial constraints. *Discrete and Computational Geometry*, 30(4):607–622, 2003.

[10] F. Geerts and B. Kuijpers. Expressing topological connectivity of spatial databases. In R.C.H. Connor and A.O. Mendelzon, editors, *Research Issues in Structured and Semistructured Database Programming*, volume 1949 of *Lecture Notes in Computer Science*, pages 224–238. Springer, 2000.

[11] F. Geerts, B. Kuijpers, and J. Van den Bussche. Linearization and completeness results for terminating transitive closure queries on spatial databases. *SIAM Journal on Computing*. To appear.

[12] F. Geerts, L. Smits, and J. Van den Bussche. N- versus (N+1)-dimensional connectivity testing of first-order queries to semi-algebraic sets. *Acta Informatica*. To appear.

[13] C. Giannella and D. Van Gucht. Adding a path connectedness operator to FO+poly. *Acta Informatica*, 38(9):621–648, 2002.

[14] M. Grohe and L. Segoufin. On first-order topological queries. *ACM Transactions on Computational Logic*, 3(3):336–358, 2002.

[15] S. Grumbach and J. Su. Queries with arithmetical constraints. *Theoretical Computer Science*, 173(1):151–181, 1997.

[16] J. Heintz and B. Kuijpers. Constraint databases, data structures, and efficient query evaluation. In B. Kuijpers and P. Revesz, editors, *Constraint Databases—Proceedings CDB 2004*, volume 3074 of *Lecture Notes in Computer Science*, pages 1–24. Springer, 2004.

[17] P.C. Kanellakis, G.M. Kuper, and P.Z. Revesz. Constraint query languages. *Journal of Computer and System Sciences*, 51(1):26–52, August 1995.

[18] S. Kreutzer. Fixed-point query languages for linear constraint databases. In *Proceedings 19th ACM Symposium on Principles of Database Systems*, pages 116–125. ACM Press, 2000.

[19] S. Kreutzer. Operational semantics for fixed-point logics on constraint databases. In R. Nieuwenhuis and A. Voronkov, editors, *Logic for Programming, Artificial Intelligence, and Reasoning—Proceedings LPAR 2001*, volume 2250 of *Lecture Notes in Computer Science*, pages 470–484, 2001.

[20] B. Kuijpers, J. Paredaens, and J. Van den Bussche. On topological elementary equivalence of closed semi-algebraic sets in the plane. *Journal of Symbolic Logic*, 65(4):1530–1555, 2000.

[21] G. Kuper, L. Libkin, and J. Paredaens, editors. *Constraint Databases*. Springer, 2000.

[22] L. Libkin. *Elements of Finite Model Theory*. Springer, 2004.

[23] C.H. Papadimitriou, D. Suciu, and V. Vianu. Topological queries in spatial databases. *Journal of Computer and System Sciences*, 58(1):29–53, 1999.

[24] P. Rigaux, M. Scholl, and A. Voisard. *Spatial Databases*. Morgan Kaufmann, 2001.

[25] J.T. Schwartz, M. Sharir, and J. Hopcroft, editors. *Planning, Geometry, and Complexity of Robot Motion*. Ablex Publishing Corporation, Norwood, New Jersey, 1987.

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

[27] A. Tarski. *A Decision Method for Elementary Algebra and Geometry*. University of California Press, 1951.

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

[29] M. Worboys and M. Duckham. *GIS: A Computing Perspective*. CRC Press, 2004.