RWTH Aachen

Department of Computer Science
Technical Report

An Existential Locality Theorem

Martin Grohe and Stefan Wöhrle

ISSN 0935–3232 · Aachener Informatik Berichte · AIB-2001-07

RWTH Aachen · Department of Computer Science · June 2001

An Existential Locality Theorem*

Martin Grohe[1] and Stefan Wöhrle[2]

[1] Laboratory for Foundations of Computer Science, University of Edinburgh,
Edinburgh EH9 3JZ, Scotland, UK.
grohe@dcs.ed.ac.uk

[2] Lehrstuhl für Informatik VII, RWTH Aachen, 52056 Aachen, Germany.
woehrle@informatik.rwth-aachen.de

Abstract. We prove an existential version of Gaifman's locality theorem and show how it can be applied algorithmically to evaluate existential first-order sentences in finite structures.

1 Introduction

Gaifman's locality theorem [12] states that every first-order sentence is equivalent to a Boolean combination of sentences saying: There exist elements $a_1, \ldots, a_k$ that are far apart from one another, and each a_i satisfies some local condition described by a first-order formula whose quantifiers only range over a fixed-size neighborhood of an element of a structure. We prove that every *existential* first-order sentence is equivalent to a *positive* Boolean combination of sentences saying: There exist elements $a_1, \ldots, a_k$ that are far apart from one another, and each a_i satisfies some local condition described by an *existential* first-order formula.

The locality of first-order logic can be explored to prove that certain properties of finite structures are not expressible in first-order logic, and it seems that this was Gaifman's main motivation. More recently, Libkin and others considered this technique of proving inexpressibility results using locality in a complexity theoretic context (see, e.g., [5, 14, 13, 15]).

A completely different application of Gaifman's theorem has been proposed in [11]: It can be used to evaluate first-order sentences in certain finite structures quite efficiently. In general, it takes time $n^{\Theta(l)}$ to decide whether a structure of size n satisfies a first-order sentence of size l, and under complexity theoretic assumptions, it can be proved that no real improvement is possible: The problem of deciding whether a given structure satisfies a given first-order sentence is PSPACE-complete [17, 19], and if parameterized by the size of the input sentence, it is complete for the parameterized complexity class AW$[*]$ [7]. The latter result implies that it is unlikely that the problem is fixed-parameter tractable (cf. [6]), i.e., that it can be solved in time $f(l) \cdot n^c$, for a function f and a constant c.

Gaifman's theorem reduces the question of whether a first-order sentence holds in a structure to the question of whether the structure contains elements that are far apart

* A conference version of this paper appears in *Proceedings of the 10th Annual Conference of the European Association for Computer Science Logic(CSL '01).*

from one another and satisfy some local condition expressed by a first-order formula. In certain structures, it is much easier to decide whether an element satisfies a local first-order formula than to decide whether the whole structure satisfies a first-order sentence. An example are graphs of bounded degree: Local neighborhoods of vertices in such graphs have a size bounded by a constant only depending on the radius of the neighborhoods, so the time needed to check whether a vertex satisfies a local condition does not depend on the size of the graph. Another, less obvious example are planar graphs. To evaluate local conditions in planar graphs, we can exploit the fact that in planar graphs neighborhoods of fixed radius have bounded tree-width [16]. In general, such a locality based approach to evaluating first-order sentences in finite structures works for classes of structures that have a property called bounded local tree-width; the class of planar graphs and all classes of structures of bounded degree are examples of classes having this property. It has been proved in [11] that for each class C of structures of bounded local tree-width there is an algorithm that, given a structure $\mathcal{A} \in$ C and a first-order sentence φ, decides whether $\mathcal{A}$ satisfies φ in time near linear in the size of the structure $\mathcal{A}$ (the precise statement is Theorem 7).

While a linear dependence on the size of the input structure is optimal, the dependence of these algorithms on the size of the input sentence leaves a lot to be desired: There is not even an elementary upper bound for the runtime in terms of the size of the sentence. Although the dependence of the algorithm on the structure size matters much more than the dependence on the size of the sentence, because usually we are evaluating small sentences in large structures,[1] it would be desirable to have a dependence on the size of the sentence that is not worse than exponential. Of course, since we are dealing with a PSPACE complete problem, we cannot expect the runtime of an algorithm to be polynomial in both the size of the input structure and the size of the input sentence.

We have observed that one of the main factors contributing to the enormous runtime of the locality based algorithms in terms of the formulas size is the number of quantifier alternations in the formula. This has motivated the present paper. We can use a variant of our existential locality theorem to improve the algorithms described above to algorithms whose runtime "only" depends doubly exponentially on the size of the input sentence.

In this paper we concentrate on the proof of our existential locality theorem, which is surprisingly complicated. This proof is presented in Section 3. The algorithmic application is outlined in Section 4.

2 Preliminaries

A *vocabulary* is a finite set of relation symbols. Associated with every relation symbol R is a positive integer called the *arity* of R. In the following, τ always denotes a vocabulary.

A τ*-structure* $\mathcal{A}$ consists of a non-empty set A, called the *universe* of $\mathcal{A}$, and a relation $R^{\mathcal{A}} \subseteq A^r$ for each r-ary relation symbol $R \in \tau$. For instance, we consider *graphs* as $\{E\}$-structures $\mathcal{G} = (G, E^{\mathcal{G}})$, where the binary relation $E^{\mathcal{G}}$ is symmetric and

[1] The generic example is the problem of evaluating SQL database queries against finite relational databases, which can be modeled by the problem of evaluating first-order sentences in finite structures.

anti-reflexive (i.e. graphs are undirected and loop-free). If $\mathcal{A}$ is a τ-structure and $B \subseteq A$, then $\langle B \rangle^{\mathcal{A}}$ denotes the substructure induced by $\mathcal{A}$ on B, that is, the τ-structure $\mathcal{B}$ with universe B and $R^{\mathcal{B}} := R^{\mathcal{A}} \cap B^r$ for every r-ary $R \in \tau$.

The formulas of *first-order logic* are build up from *atomic formulas* using the usual Boolean connectives and existential and universal quantification over the elements of the universe of a structure. Remember that an *atomic formula*, or *atom*, is a formula of the form $x = y$ or $R(x_1, \dots, x_r)$, where R is an r-ary relation symbol. The set of all variables of a formula φ is denoted by $\text{var}(\varphi)$. A *free variable* in a first-order formula is a variable x not in the scope of a quantifier $\exists x$ or $\forall x$. The set of all free variables of a formula φ is denoted by $\text{free}(\varphi)$. A *sentence* is a formula without free variables. The notation $\varphi(x_1, \dots, x_k)$ indicates that all free variables of the formula φ are among $x_1, \dots, x_k$; it does not necessarily mean that the variables $x_1, \dots, x_k$ all appear in φ. For a formula $\varphi(x_1, \dots, x_k)$, a structure $\mathcal{A}$, and $a_1, \dots, a_k \in A$ we write $\mathcal{A} \models \varphi(a_1, \dots, a_k)$ to say that $\mathcal{A}$ satisfies φ if the variables $x_1, \dots, x_k$ are interpreted by the vertices $a_1, \dots, a_k$, respectively.

The *weight* of a first-order formula φ is the number of quantifiers $\exists x$ and $\forall x$ occurring in φ.

A first-order formula is *existential* if it contains no universal quantifiers and if every existential quantifier occurs in the scope of an even number of negation symbols. A *literal* is an atom or a negated atom. A *conjunctive query with negation* is a formula of the form $\exists \bar{x} \bigwedge_{i=1}^{m} \lambda_i$, where each λ_i is a literal. Every existential formula φ of weight w and length l is equivalent to a disjunction of at most 2^l conjunctive queries with negation, each of which is of weight at most w and length at most l.

We often denote tuples $a_1 \dots a_k$ of elements of a set A by $\bar{a}$, and we write $\bar{a} \in A$ instead of $\bar{a} \in A^k$. Similarly, we denote tuples of variables by $\bar{x}$.

Our underlying model of computation is the standard RAM-model with addition and subtraction as arithmetic operations (cf. [1, 18]). In our complexity analysis we use the uniform cost measure. Structures are represented on a RAM in a straightforward way by listing all elements of the universe and then all tuples in the relations. For details we refer the reader to [10]. We define the *size* of a τ-structure $\mathcal{A}$ to be $||\mathcal{A}|| := |A| + \sum_{R \in \tau \ r\text{-ary}} r \cdot |R^{\mathcal{A}}|$; this is the length of a reasonable representation of $\mathcal{A}$ (if we suppress details that are inessential for us). We fix some reasonable encoding for first-order formulas and denote by $||\varphi||$ the size of the encoding of a formula φ.

2.1 Gaifman's Locality Theorem

The *Gaifman graph* of a τ-structure $\mathcal{A}$ is the graph $\mathcal{G}_{\mathcal{A}}$ with vertex set A and an edge between two vertices $a, b \in A$ if there exists an $R \in \tau$ and a tuple $a_1 \dots a_k \in R^{\mathcal{A}}$ such that $a, b \in \{a_1, \dots, a_k\}$. The *distance* $d^{\mathcal{A}}(a, b)$ between two elements $a, b \in A$ of a structure $\mathcal{A}$ is the length of the shortest path in $\mathcal{G}_{\mathcal{A}}$ connecting a and b. For $r \geq 1$ and $a \in A$, we define the *r-neighborhood* of a in $\mathcal{A}$ to be $N_r^{\mathcal{A}}(a) := \{b \in A \mid d^{\mathcal{A}}(a, b) \leq r\}$. For a subset $B \subseteq A$ we let $N_r^{\mathcal{A}}(B) := \bigcup_{b \in B} N_r^{\mathcal{A}}(b)$.

For every $r \geq 0$ there is an existential first-order formula $\delta_r(x, y)$ such that for all τ-structures $\mathcal{A}$ and $a, b \in A$ we have $\mathcal{A} \models \delta_r(a, b)$ if, and only if, $d^{\mathcal{A}}(a, b) \leq r$.