

Declarative PTIME Queries to Relational Databases

Patrick Doherty and Witold Łukaszewicz and Andrzej Szalas

Abstract

In this paper, we consider the problem of expressing and computing PTIME queries to relational deductive databases in a purely declarative query language, SHQL (Semi-Horn Query Language). Assuming the relational databases in question are ordered, we show that all SHQL queries are computable in PTIME and the whole class of PTIME queries is expressible in SHQL. Although similar results have been proven for fixpoint languages and extensions to datalog, the claim is that SHQL has the advantage of being purely declarative, where the negation operator is interpreted as classical negation, mixed quantifiers may be used and a query is simply a restricted first-order theory not limited by the rule-based syntactic restrictions associated with logic programs in general. We describe the PTIME algorithm used to compute queries in SHQL which is based in part on quantifier elimination techniques and also extend the method to incomplete relational databases using intuitions from Circumscription.

The first author has been supported in part by the Swedish Research Council for the Engineering Sciences (TFR). The second and third authors are affiliated with the Institute of Informatics, Warsaw University, and have been supported in part by KBN grant 3 P406 019 06.

Declarative PTIME Queries to Relational Databases

Patrick Doherty* Witold Łukaszewicz† Andrzej Szalas†

Abstract

In this paper, we consider the problem of expressing and computing PTIME queries to relational deductive databases in a purely declarative query language, SHQL (Semi-Horn Query Language). Assuming the relational databases in question are ordered, we show that all SHQL queries are computable in PTIME and the whole class of PTIME queries is expressible in SHQL. Although similar results have been proven for fixpoint languages and extensions to datalog, the claim is that SHQL has the advantage of being purely declarative, where the negation operator is interpreted as classical negation, mixed quantifiers may be used and a query is simply a restricted first-order theory not limited by the rule-based syntactic restrictions associated with logic programs in general. We describe the PTIME algorithm used to compute queries in SHQL which is based in part on quantifier elimination techniques and also extend the method to incomplete relational databases using intuitions from Circumscription.

1 Introduction

In this paper, we consider the problem of expressing and computing PTIME queries to relational deductive databases in a purely declarative query language, SHQL (Semi-Horn Query Language). Assuming the relational databases in question are ordered, we show that all SHQL queries are computable in PTIME and the whole class of PTIME queries is expressible in SHQL. Similar results have been proven for fixpoint languages and extensions to datalog, but the claim will be that these languages are not purely declarative and are far from providing a natural means of expressing queries.

Much recent activity in the area of deductive databases has focused on the toy language *datalog* and its extensions which integrate recursion with negation. When adding negation to datalog, this requires defining a semantics for negative facts. There are many choices as to such a semantics and these choices influence not only the natural interpretation of the negation symbol in a query, but the expressiveness of the language. For example, stratified semantics requires syntactic restrictions on the use of negation in a datalog⁷ program, while well-founded semantics, although not requiring syntactic restrictions, does use a 3-valued semantics to interpret the meaning of a program. In addition, while well-founded semantics is equivalent to the fixpoint queries, stratified semantics is strictly weaker.

An important aspect of query language design is to achieve a good balance between the expressiveness of the language and the complexity of evaluating queries in the language. In addition to expressiveness and efficiency, the language should be natural to use. Although it can be argued that extended datalog languages achieve the goals of expressiveness and efficiency in theory, one can debate the naturalness of using datalog as a query language. For instance, the variations in interpretation already discussed can be quite confusing for a normal user of the query language.

*Department of Computer and Information Science, Linköping University, S-581 83 Linköping, Sweden, e-mail: patdo@ida.liu.se. Supported in part by the Swedish Research Council for Engineering Sciences (TFR).

†Institute of Informatics, Warsaw University, ul. Banacha 2, 02-097 Warsaw, Poland, e-mail: witlu,szalas@mimuw.edu.pl. Supported in part by KBN grant 3 P406 019 06.

The procedural leakage into the language which is a result of alternative non-classical interpretations of the negation operator and the various syntactic restrictions placed on datalog programs, is not only unnatural but tends to violate the basic tenets of a good declarative query language.

On the other hand, SHQL is a purely declarative query language. Use of negation in a query is interpreted as classical negation, a class of mixed quantifiers is allowed in queries, and intentional and extensional predicates may occur anywhere in the query. SHQL is not rule-based and a query is expressed as a theory consisting of semi-Horn formulas.

SHQL is used as follows. Given the task of computing a definition of an intensional predicate Q (or asking whether a tuple is an instance of Q) relative to a relational database B consisting of the relations R_1, \dots, R_n , we first provide an implicit definition of Q in terms of a SHQL theory, $\Theta(Q)$, which is essentially a conjunction of semi-Horn formulas using any of R_1, \dots, R_n , and Q . The theory $\Theta(Q)$ is only constrained by the fact that it must be semi-Horn. All quantifiers and logical connectives are interpreted classically. The goal is to compute an explicit definition of Q in PTIME.

The computation process can be described in two stages. In the first stage, we provide a PTIME compilation process which uses a quantifier elimination algorithm called the DLS algorithm [3]. An extension for fixpoint formulas is called the G-DLS algorithm [4, 5]. The DLS algorithm takes as input a second-order formula and returns a logically equivalent first-order formula, or terminates with failure, where failure does not mean there is not a reduction, but simply that the algorithm can not find one. The G-DLS algorithm is a generalization of the DLS algorithm and returns logically equivalent fixpoint formulas for a wider class of input. Both algorithms can be combined into one algorithm which we denote by DLS* (see [5]). Given the SHQL query, $\Theta(Q)$, we prefix it with an existential quantifier and input the formula $\exists Q.\Theta(Q)$ to DLS*. If the query is first-order definable then the output will be a logically equivalent first-order formula expressing an explicit definition of Q . The output is computed in PTIME and LOGSPACE (in the size of the database). If the query is not first-order definable, then the output will be a logically equivalent fixpoint formula expressing an explicit definition of Q . In this case, output is computed in PTIME. Note that this technique can be used for theories outside the semi-Horn class, but neither the complexity results nor a successful reduction are guaranteed.

In the second stage, we use the explicit definition of Q output in the first stage to compute a suitable relation in the relational database that satisfies Q . Before computing the output relation, we first check to see that such a relation exists relative to the database. Suppose $\Theta(Q)$ is the original query, B the relational database and $\Theta'(Q)$ the output of DLS* given the input $\exists Q.\Theta(Q)$. We say that the query $\Theta(Q)$ is a *coherent query relative to B* if $B \models \Theta'(Q)$. Assuming this is the case, we know that the output relation exists and can now compute the answer. Both checking that the query is coherent ($B \models \Theta'(Q)$) and computing the output relation can be done efficiently because calculating fixpoint queries and fixpoint satisfiability checking over finite domains are both in PTIME (see Immerman [7], Sazonov [12], Vardi [14]).

Note that although the combined problem of finding out whether an implicit query $\Theta(Q)$ to a database exists, checking that the query is coherent, and explicitly computing the answer is in general NP-complete (in the size of the database), as was shown by Fagin (see Immerman [8]), our method which applies quantifier elimination techniques to semi-Horn theories makes the problem solvable in polynomial time for this special case. Most importantly, SHQL is a highly expressive language which covers all PTIME queries and is at the same time purely declarative. Querying with SHQL is as natural as querying with classical logic and the compilation step is completely transparent to the user.

The rest of the paper will be structured as follows. In Section 2, we provide preliminary definitions. In Section 3, we describe Ackermann's Lemma and the Fixpoint Theorem, which provide the formal basis for the DLS* algorithm. In fact, using the full algorithm is not necessary in order to achieve our goals. Some more direct syntactic manipulations together with these theorems provably achieve the same reduction results and are discussed in Section 4.2. In Section 4, we provide a detailed

description of our two stage method and discuss a technique which permits queries with more than one intensional predicate. In Section 5, we provide a number of examples which demonstrate the naturalness of SHQL and the proposed querying method. Finally, in Section 6, we extend the method to relational databases with incomplete information. We then conclude with a discussion.

2 Definitions

In what follows, by a *theory* we always mean a finite set of axioms. Thus theories can be transformed into formulas (conjunctions of axioms).

Definition 2.1 A *relational database* B , is a first order structure $\langle U, r_1^{a_1}, \dots, r_k^{a_k}, c_1, \dots, c_l \rangle$, where

- U is a finite set,
- for $1 \leq i \leq k$, $r_i^{a_i}$ is an a_i -ary relation on U , i.e. $r_i^{a_i} \subseteq U^{a_i}$, and
- $c_1, \dots, c_l \in U$ are constants.

By a *signature* of B we mean a signature containing relation symbols $R_1^{a_1}, \dots, R_k^{a_k}$ and constant symbols C_1, \dots, C_l together with equality $=$. ■

According to accepted terminology in the literature (introduced in [11]), a deductive database consists of two parts: an extensional and intensional database. The extensional database is usually equivalent to a traditional relational database and the intensional database contains a set of definitions of relations that are not explicitly stored in the database. Accordingly, we have the following definition.

Definition 2.2 By a *deductive database* we understand a relational database augmented with an additional set of formulas defining fresh relations in terms of a chosen logic. The relational database is called *extensional* and the set of formulas is called an *intensional database*. We say that a relation (relation symbol) is *intensional* in a database if it appears in the intensional database only, otherwise it is called *extensional*. ■

Definition 2.3 We say that a formula Φ is *positive* w.r.t. a predicate P iff P appears under no negation sign in Φ (in negation normal form). Dually, we say that Φ is *negative* w.r.t. P iff all occurrences of P have the form $\neg P$ and $\neg P$ appears under no negation sign in Φ . ■

Definition 2.4 Let $\mathcal{L}_{\mathcal{I}}$ be the classical first-order logic and σ be a signature. By the *fixpoint calculus* over σ , denoted by $\mathcal{L}_{\mathcal{F}}^{\sigma}$, we understand this to be the logic obtained from $\mathcal{L}_{\mathcal{I}}$ by extending it with the least fixpoint operator $\mu P.\Phi(P)$, where Φ is positive w.r.t. P . We define $\nu P.\Phi(P)$ as $\neg\mu\neg P.\neg\Phi(P)$ ¹. ■

Note that $\mu P(\bar{x}).\Phi(P)$ is the least (w.r.t. implication) formula $\Psi(\bar{x})$ such that

$$\Psi(\bar{x}) \equiv \Phi(P \leftarrow \Psi(\bar{x})).$$

Every formula $\Phi(P)$ which is positive w.r.t. P is monotone and therefore, by the Knaster & Tarski fixpoint theorem, the fixpoints we consider are well defined.

Definition 2.5 Let $B = \langle U, r_1^{a_1}, \dots, r_k^{a_k}, c_1, \dots, c_l \rangle$ be a relational database and let σ be a signature of B .

- By a *first-order query language for* B , denoted by $\mathcal{L}_{\mathcal{I}}^B$ we mean the classical first-order logic over signature σ .

¹ Observe that a more readable form of this formula can be obtained by replacing P by $\neg P$ and writing $\neg\mu P.\neg\Phi(P)$

- By a *fixpoint query language for B*, denoted by $\mathcal{L}_{\mathcal{F}}^B$ we mean the fixpoint calculus over signature σ .
- By an *implicit query to B* we mean a classical first-order formula $\Theta(Q)$ over signature σ augmented with an additional relation symbol Q (representing the relation to be calculated).

■

Observe that in the case of an implicit query, say $\Theta(Q)$, it is natural to demand that Q represents the minimal or maximal relation s satisfying Θ , provided that such s exists². Let us note that maximizing a relation corresponds to minimizing its complement³. Accordingly, and without loss of generality, we shall focus on minimizing relations.

Definition 2.6 Let $B = \langle U, r_1^{a_1}, \dots, r_k^{a_k}, c_1, \dots, c_l \rangle$ be a relational data base.

- The semantics of the query language \mathcal{L}_T^B is defined as in the case of first-order logic, assuming that for $1 \leq i \leq k$, relation symbols $R_i^{a_i}$ are interpreted as relations $r_i^{a_i}$ and for $1 \leq j \leq l$, constant symbols C_j are interpreted as constants c_j .
- The semantics of the query language $\mathcal{L}_{\mathcal{F}}^B$ is defined by extending the definition of the semantics of \mathcal{L}_T^B , assuming that $\mu Q. \Phi(Q)$ represents the least (w.r.t. \subseteq) relation s such that $B \models Q = \Phi(Q)$ with Q interpreted as s .
- The semantics of an implicit query $\Theta(Q)$ is defined as the least (w.r.t. \subseteq) relation s such that $B \models \Theta(Q)$ with Q interpreted as s , provided that such s exists. A formula expressing the existence of s is called the *coherence condition* for $\Theta(Q)$. ■

It is often convenient to reference some “columns” of a relation in a database. We often do this by extending the signature with relation symbols corresponding to columns. For example, given a relation $person \subseteq String \times Integer$, we might want to refer to the first column using the predicate symbol $Name$ where $Name(x, john)$ means that the name of x is $john$. If a column, being represented by, say C , contains boolean values then we write $C(x)$ to mean $C(x, \top)$ and $\neg C(x)$ to mean $C(x, \perp)$. This notation for referencing columns will prove to be useful in Section 6, where we consider incomplete databases.

3 Ackermann’s Lemma and a Fixpoint Theorem

In the introduction, we stated that the DLS* algorithm was a combination of two separate algorithms, DLS and G-DLS. DLS* works as follows. The input to DLS* is first passed to the DLS algorithm. If the input can be put into what we call a *Ackermann reducible formula* than DLS* outputs a logically equivalent first-order formula. If the DLS algorithm terminates with failure, then the input is passed to the G-DLS algorithm. If the input can be put into what we call a *semi-Horn formula* then the DLS* algorithm outputs a fixpoint formula logically equivalent to the input. Of course, certain optimizations can be made which combine the two algorithms in a more efficient manner. The basis for both the DLS and G-DLS algorithms are two theorems which we describe below. Both provide a means of eliminating quantifiers which bind predicate variables. These two theorems provide the formal basis for the compilation step described previously and called stage one, where an SHQL query is first prefixed with an existential quantifier which binds the intensional predicate whose explicit definition we would like to generate and compute. The second-order query is then passed to the DLS* algorithm.

The following lemma was proved by Ackermann in [1] (for an alternative proof see also [13]).

Lemma 3.1 Let P be a predicate variable and $\Phi(\bar{x}, \bar{z})$, $\Psi(P)$ be formulas without second-order quantification. Let Φ contain no occurrences of P at all. Then the following equivalences hold:

²The existence of s means that $\Theta(Q)$ is consistent with the database.

³This does not have to hold when a database is allowed to contain incomplete information - see Section 6

Let $\Psi(\neg P)$ be negative w.r.t. P , then

$$\exists P \forall \bar{x} [P(\bar{x}) \vee \Phi(\bar{x}, \bar{z})] \wedge \Psi(\neg P) \equiv \Psi(\neg P \leftarrow \Phi(\bar{x}, \bar{z})) \quad (1)$$

Let $\Psi(P)$ be positive w.r.t. P , then

$$\exists P \forall \bar{x} [\neg P(\bar{x}) \vee \Phi(\bar{x}, \bar{z})] \wedge \Psi(P) \equiv \Psi(P \leftarrow \Phi(\bar{x}, \bar{z})), \quad (2)$$

where in the right-hand formulae the arguments \bar{x} of Φ are to be substituted by the respective actual arguments of P (renaming the bound variables whenever necessary). ■

The following theorem, extending Lemma 3.1, is proved in [10].

Theorem 3.1 Assume that all occurrences of the predicate symbol P in the formula Ψ have only variables as arguments. Then the following equivalences hold:

Let $\Phi(\neg P)$ and $\Psi(\neg P)$ be negative w.r.t. P , then

$$\exists P \forall \bar{y} [P(\bar{y}) \vee \Phi(\neg P)] \wedge [\Psi(\neg P)] \equiv \Psi[\neg P \leftarrow \nu \neg P(\bar{y}).\Phi(\neg P)], \quad (3)$$

Let $\Phi(P)$ and $\Psi(P)$ be positive w.r.t. P , then

$$\exists P \forall \bar{y} [\neg P(\bar{y}) \vee \Phi(P)] \wedge [\Psi(P)] \equiv \Psi[P \leftarrow \nu P(\bar{y}).\Phi(P)], \quad (4)$$

where the above substitutions exchange the variables bound by fixpoint operators by the corresponding actual variables of the substituted predicate. ■

Formula (2) of Lemma 3.1 and formula (4) of Theorem 3.1 are applied in the case of minimizing relations, while formulas (1) and (3) are applied when maximizing relations.

Lemma 3.1 is subsumed by Theorem 3.1. Moreover, any fixpoint formula of the form $\nu P.\Psi$, where Ψ does not contain P , is equivalent to Ψ . Thus one can, in all cases, use Theorem 3.1 and simplify the resulting formulas by applying this equivalence. This optimization diverges from the conceptual description we have been using when describing the DLS* algorithm, but it will simplify the detailed description of the query method described in the next section.

4 The Method

We first observe that the problem whether a result of an implicit query $\Theta(Q)$ to a database B exists reduces to the question whether the second-order formula $\exists Q \Theta(Q)$ is satisfied in B . By Fagin's theorem the problem is *NP*-complete in the size of B (see [8]). In what follows we concentrate on selecting a class of implicit queries for which the problem is in *PTIME*.

Conceptually, the SHQL query method consists of four steps:

1. State a query $\Theta(Q)$ to a relational database B in SHQL, where $\Theta(Q)$ is a semi-Horn formula (for the definition of semi-Horn formulas see Section 4.1). Prefix the query with an existential quantifier binding the intensional predicate whose explicit definition in terms of $\Theta(Q)$ we would like to make explicit. The input to the compilation stage is $\exists Q.\Theta(Q)$.
2. Pass the input $\exists Q.\Theta(Q)$ to the DLS* algorithm. Assuming the input is semi-Horn, the algorithm will return either a logically equivalent first-order formula or a fixpoint formula. Call the output $\Theta'(Q)$.
3. Before explicitly computing the answer to the original query $\Theta(Q)$, check to make sure the query is *coherent* relative to B . We do this by essentially checking that $\Theta'(Q)$ is satisfied by B .

4. If the query is coherent, than compute the definition of Q (or check whether a tuple belongs to Q).

Provided the input is semi-Horn, all steps in the method can be computed in PTIME. In the following subsections, we will formally define the query language, describe and justify each of steps 2-4 with appropriate theorems, and conclude with a representation theorem characterizing the expressiveness and descriptive complexity of the query language.

4.1 The Semi-Horn Query Language (SHQL)

In previous sections, we discussed SHQL informally stating that queries to the database had to be in what we called *semi-Horn* form. In practice, we can do more. A query $\Theta(Q)$ can be any formula in a classical first-order language. If it is, or can be transformed into semi-Horn form, then the DLS* algorithm is guaranteed to generate an explicit definition of Q which is logically equivalent to $\exists Q.\Theta(Q)$. If it is not, the DLS* algorithm may still terminate successfully and steps 2-4 above would still apply. In this section, we will make these intuitions precise.

We shall consider two types of formulas of the form

$$\Phi_1(Q) \wedge \Phi_2(Q), \quad (5)$$

where $\Phi_2(Q)$ is any first-order formula negative w.r.t. Q . We call these two types Ackermann-reducible formulas and semi-Horn formulas. They are defined as follows:

- *Ackermann-reducible formulas* (w.r.t. Q) are of the form (5) for which $\Phi_1(Q)$ is a conjunction of formulas of the form $\forall \bar{x}(Q(\bar{t}) \vee \Psi)$, where Ψ is an arbitrary Q -free first-order formula
- *semi-Horn formulas* (w.r.t. Q) are of the form (5) for which $\Phi_1(Q)$ is a conjunction of formulas of the form $\forall \bar{x}(Q(\bar{t}) \vee \Psi(\neg Q))$, and Ψ is an arbitrary first-order formula negative w.r.t. Q .

The negative dual forms are obtained by substituting Q by $\neg Q$ in the definitions, making Ψ an arbitrary first-order formula positive w.r.t. Q , and making $\Phi_2(Q)$ negative w.r.t. Q . In this case we are able to find the greatest solution for $\neg Q$, that is, a minimal solution for Q .

Ackermann-reducible formulas are also semi-Horn formulas. However, it is important to isolate this class of formulas because these define first-order expressible queries. If the initial query $\Theta(Q)$ can be transformed into this form, then in the compilation step where $\exists Q.\Theta(Q)$ is given as input to the DLS* algorithm, the call to DLS will return a logically equivalent first-order formula. In fact, the DLS* algorithm, when successful, basically transforms $\Theta(Q)$ into one of the above forms (or their negative duals). Consequently, it is important to note that the method can be made more general by generating solutions for arbitrary formulas $\Gamma(Q)$ which although not semi-Horn, are reducible by the DLS* algorithm.

Note that any conjunction of semi-Horn formulas w.r.t. Q can be transformed into the following form:

$$\forall \bar{x}[\Phi(\bar{x}, \bar{z}_i, Q) \supset Q(\bar{x})] \wedge \Psi(\neg Q), \quad (6)$$

where $\Psi(\neg Q)$ is an arbitrary first-order formula negative w.r.t. Q and $\Phi(\bar{x}, \bar{z}_i, Q)$ is positive w.r.t. Q , or its dual,

$$\forall \bar{x}[\Phi(\bar{x}, \bar{z}_i, \neg Q) \supset \neg Q(\bar{x})] \wedge \Psi(Q), \quad (7)$$

where $\Psi(Q)$ is an arbitrary first-order formula positive w.r.t. Q and $\Phi(\bar{x}, \bar{z}_i, \neg Q)$ is negative w.r.t. Q . To see this it suffices to use the following equivalence that allows us to combine two semi-Horn formulas into a single semi-Horn formula:

$$\begin{aligned} & [\forall \bar{x}(\Phi(\bar{x}, \bar{z}_i, Q) \supset Q(\bar{x})) \wedge \Psi(\neg Q)] \wedge [\forall \bar{x}(\Phi'(\bar{x}, \bar{z}'_i, Q) \supset Q(\bar{x})) \wedge \Psi'(\neg Q)] \\ & \equiv [\forall \bar{x}((\Phi(\bar{x}, \bar{z}_i, Q) \vee \Phi'(\bar{x}, \bar{z}'_i, Q)) \supset Q(\bar{x})) \wedge (\Psi(\neg Q) \wedge \Psi'(\neg Q))] \end{aligned} \quad (8)$$

A similar equivalence applies for the dual forms. Moreover, if we have many intensional predicates, we can easily encode these by a single predicate which has a vector of additional boolean variables distinguishing between various relations. For instance, if we have two predicates, say $P(\bar{x})$ and $Q(\bar{y})$, then we can use a single predicate $R(z, \bar{x}, \bar{y})$ such that $R(\top, \bar{x}, \bar{y})$ means $P(\bar{x})$ and $R(\perp, \bar{x}, \bar{y})$ means $Q(\bar{y})$. The resulting formulas are still semi-Horn formulas. Thus, we can safely assume that we always deal with a single intensional predicate in our queries.

It is worth emphasizing here that semi-Horn formulas are strictly more expressive than Horn clauses. For instance, semi-Horn formulas express the complement of a relation which is not expressible by Horn clauses (see e.g. [2]). Also existential quantifiers in the scope of universal quantifiers are not, in general, reducible to Horn clauses, but are allowed in semi-Horn formulas.

Let us now introduce the definition of declarative queries and declarative query language. As we shall see in Theorem 4.3 all the queries of the language are computable in polynomial time. Moreover, the whole class of PTIME queries is covered by the language.

Definition 4.1 By a *declarative query* we mean any implicit query expressed as a semi-Horn formula. By a *declarative query language* we mean a first-order query language augmented with declarative queries, assuming that the underlying signature contains a relation that, on the semantic side, linearly orders domains of databases. ■

4.2 The Compilation, Coherence, and Computing Steps

It is easily observed that Ackermann-reducible formulas are reducible to classical first-order formulas by applying the DLS algorithm which is based on Lemma 3.1 and that semi-Horn formulas are reducible to fixpoint formulas by applying the G-DLS algorithm which is based on Theorem 3.1 (see e.g. [3, 4]).

We also observe two additional facts concerning Lemma 3.1 and Theorem 3.1. Namely, assume we are given a formula $\exists Q\Phi(Q)$. By the proofs of Lemma 3.1 and Theorem 3.1, where reduction is successful (which is always the case for semi-Horn queries) one gets:

- a first-order (or fixpoint) definition of Q (this definition is used in suitable substitutions in the resulting formulas), and
- a first-order (or fixpoint) formula equivalent to the input formula.

The first observation justifies the generation of a fixpoint or first-order formula that explicitly defines the query. The second observation justifies the generation of the coherence condition for a query expressed as a fixpoint or first-order formula. Note that the explicit definition and coherence condition are either both first-order or both fixpoint formulas. The coherence condition allows us to check whether the output relation exists. If we know that the output relation exists, we can calculate the answer using the formula obtained via the first observation. Both the coherence check and calculation of the output relation can be done in polynomial-time by using an algorithm for calculating fixpoint queries and for checking fixpoint satisfiability over finite domains which is described in [7]. The process of calculating the output relation and checking coherence can be optimized by noting that more or less the same fixpoint formula appears in the explicit definition and the coherence condition.

The following theorem easily follows from a result in [4].

Theorem 4.1

For any formula $\Theta(Q)$ of the form (6):

- the explicit definition of Q is given by $Q(\bar{x}) \equiv \mu Q(\bar{x}).\Phi(\bar{x}, \bar{z}_i, Q)$, and
- the coherence condition for $\Theta(Q)$ is $\Psi(Q \leftarrow \mu Q(\bar{x}).\Phi(\bar{x}, \bar{z}_i, Q))$.

For any formula $\Theta(Q)$ of the form (7):

- the explicit definition of Q is given by $Q(\bar{x}) \equiv \nu Q(\bar{x}).\neg\Phi(\bar{x}, \bar{z}_i, \neg Q)$, and
- the coherence condition for $\Theta(Q)$ is $\Psi(Q \leftarrow \nu Q(\bar{x}).\neg\Phi(\bar{x}, \bar{z}_i, \neg Q))$. ■

As a consequence we have the following theorem.

Theorem 4.2

For any formula $\Theta(Q)$ of the form (6), where Φ does not contain Q :

- the explicit definition of Q is given by $Q(\bar{x}) \equiv \Phi(\bar{x}, \bar{z}_i)$, and
- the coherence condition for $\Theta(Q)$ is $\Psi(Q \leftarrow \Phi(\bar{x}, \bar{z}_i))$.

For any formula $\Theta(Q)$ of the form (7), where Φ does not contain Q :

- the explicit definition of Q is given by $Q(\bar{x}) \equiv \neg\Phi(\bar{x}, \bar{z}_i)$, and
- the coherence condition for $\Theta(Q)$ is $\Psi(Q \leftarrow \neg\Phi(\bar{x}, \bar{z}_i))$. ■

4.3 A Representation Theorem

For the declarative query language defined in Definition 4.1 we have the following theorem.

Theorem 4.3 Let B be a relational database.

- Any implicit query $\Theta(Q)$ to B , where $\Theta(Q)$ is a semi Horn formula, is computable in polynomial time in the size of the database. Consequently, any query expressed in the declarative query language defined in Definition 4.1 is in PTIME. If $\Theta(Q)$ is an Ackermann reducible formula, then Q is computable in logarithmic space.
- Any PTIME query can be expressed in the declarative query language defined in Definition 4.1 provided that the domain of B is linearly ordered.

Proof The first part of the theorem can be proved by noticing that semi-Horn formulas reduce to fixpoint formulas and using the well-known fact that fixpoint queries are computable in polynomial time (see e.g. [2, 7]). If $\Theta(Q)$ is an Ackermann reducible formula, then both Q and the coherence condition are expressed in classical first-order logic and therefore are computable in logarithmic space.

To prove the second part of the theorem we use the following known facts:

- if the domain of the database is linearly ordered then fixpoint queries express all PTIME queries (see [2, 7, 12, 14])
- under the above assumption all fixpoint queries can be expressed by taking a single fixpoint of a first-order formula with positive occurrences of the calculated predicate (followed by classical first-order operations) (see [2, 6, 7]).

According to the above, it is sufficient to prove that a single fixpoint can be defined by an implicit query. Assume that the fixpoint to be defined is $\mu P(\bar{x}).\Psi(P)$, where Ψ is positive w.r.t. P . By theorem 4.1, the implicit query that defines this fixpoint is then simply the formula $\forall \bar{x}(\Psi(P) \supset P(\bar{x}))$. ■

5 Examples

In this section, we provide a number of examples which demonstrate both the expressiveness of semi-Horn queries and how the formal results may be applied practically.

Example 5.1 This example demonstrates how the intensional predicate Q and the extensional predicates R , S , and E may be used anywhere in the query. In particular, in comparison with rule-based queries such as logic programming or datalog, both intensional and extensional predicates may occur in both the head and body of any implication.

Assume we have a database B , containing information about whether persons are *rich*, *smart*, or *experienced*, denoted by the unary extensional predicates, R , S , and E , respectively. Suppose we are interested in selecting all rich persons and perhaps some others and we only want to consider those who are smart and experienced. Let Q denote the unary extensional predicate that describes the required relation. The first condition is then expressed by the formula

$$\forall x(R(x) \supset Q(x)),$$

while the second condition is expressed by the formula

$$\forall x(Q(x) \supset (S(x) \wedge E(x))).$$

The implicit query $\Theta(Q)$ is then defined as the conjunction of the above formulas, where we are interested in obtaining the greatest relation Q satisfying

$$\forall x(Q(x) \supset (S(x) \wedge E(x))) \wedge \forall x(R(x) \supset Q(x)).$$

After removing the implication sign, we have the equivalent,

$$\forall x(\neg Q(x) \vee (S(x) \wedge E(x))) \wedge \forall x(\neg R(x) \vee Q(x)).$$

In order to maximize Q we will minimize its negation by using the dual form (7). $\Theta(Q)$ can be rewritten as $\Theta(\neg Q)$,

$$\forall x(\neg(S(x) \wedge E(x)) \supset \neg Q(x)) \wedge \forall x(\neg R(x) \vee Q(x))$$

which is easily observed to be of the form (7), where $\Phi(\bar{x})$ is $\neg(S(x) \wedge E(x))$ and $\Psi(Q)$ is $\forall x(\neg R(x) \vee Q(x))$.

According to Theorem 4.2 the following formula of the form $\Psi(Q \leftarrow \neg\Phi(\bar{x}))$ is the suitable coherence condition:

$$\forall x(\neg R(x) \vee (S(x) \wedge E(x))).$$

Observe that our query forces this condition (by transitivity of implication), Thus, for instance, if a database contains an element e such that $R(e)$ and $\neg S(e)$, then the query is inconsistent with the database.

Now from Theorem 4.2 we obtain the explicit definition of $\neg Q$, which is $\neg Q(x) \equiv \neg(S(x) \wedge E(x))$. Consequently, the explicit definition for Q is:

$$Q(x) \equiv S(x) \wedge E(x).$$

■

Example 5.2 Consider the database containing a binary relation r . The following implicit query $\Theta(S)$ defines S as the transitive closure of r :

$$\forall x \forall y (R(x, y) \supset S(x, y)) \wedge \forall x \forall y \forall z ((S(x, y) \wedge S(y, z)) \supset S(x, z)).$$

The above formula is equivalent (and turned by the G-DLS algorithm during execution) into:

$$\exists S[\forall x \forall z (S(x, z) \vee (\neg R(x, z) \wedge \forall y (\neg S(x, y) \vee \neg S(y, z)))].$$

It is easily observed that this formula has the form (6), where $\Phi(x, z, S)$ is $\neg(\neg R(x, z) \wedge \forall y (\neg S(x, y) \vee \neg S(y, z)))$, and $\Psi(\neg S)$ is \top . Thus, according to Theorem 4.1, the coherence condition for $\Theta(S)$ is $\Psi(\neg S \leftarrow \mu S(x, z). \Phi(x, z))$. Since Ψ is \top and has no negative occurrences of S , its coherence condition is \top , which means that the required relation always exists. Moreover, we have the following explicit definition for $S(x, z)$:

$$S(x, z) \equiv \mu S(x, z). [R(x, z) \vee \exists y (S(x, y) \wedge S(y, z))].$$

■

6 Databases with Incomplete Information

So far, we have assumed that databases only contain complete information. In many AI applications, one is often confronted with the problem of using incomplete information. This could be due to the fact that we simply lack information about the properties and relations of certain objects, or that we purposely represent information incompletely for efficiency reasons. In this section, we show how one can extend the current querying method when applied to incomplete relational databases. We will only provide a proof of concept and deal with a particular class of queries leaving precise characterizations and variations of the approach for future research.

In order to simplify matters, we will assume that there is a special domain value $*$, denoting the undefined value. For example, if the database contains two pairs $\{(mary, smith), (john, *)\}$ then the first pair represents *mary smith* and the second pair represents a *john* whose family name is unknown. We will also assume that if a relation contains undefined values then it cannot be directly represented in a query. Instead, we introduce new relation symbols corresponding to its “columns”. We also assume that if a “column” can contain undefined fields, then queries are either positive or negative w.r.t. the relation symbol corresponding to the column. This is a technical assumption that allows us to proceed without any further complications, but in certain cases can be relaxed.

Of course, since we are now dealing with incomplete information, we have to make a choice regarding the semantics of partially defined predicates. We will base this choice on intuitions from AI, where second-order circumscription is used to reason about incomplete information. More precisely, we will capitalize on the partitioning of predicates in a circumscription policy into those that are minimized, fixed, or varied. When minimizing an answer (or maximizing its negation) we will allow some partially defined predicates to vary and leave some others fixed.

Let us assume that we are given an implicit query $\Theta(Q)$. In order to calculate a coherence condition and an explicit definition of Q we proceed as before. What will differ is the algorithm used for calculating the value of the coherence condition and of the relation in question. The difficulties occur when we have to calculate a value of a relation that is undefined. In such a case, we will use the following policy:

- If a relation symbol occurs positively in the query and the relation is allowed to vary then we assume \top as its value.
- If a relation symbol occurs negatively in the query and the relation is allowed to vary then we assume \perp as its value.
- if a relation is undefined for some object and the relation is fixed then we assume that the whole query is undefined for the object.

This solution precisely reflects the circumscription principle, where the possibly undefined predicates are allowed to vary in order to maximize the minimization of a predicate. The justification and correctness of this policy easily follows from the fact that both the explicit definition and the coherence condition are monotone w.r.t. predicates occurring only positively and anti-monotone w.r.t. predicates occurring only negatively.

Example 6.1 Assume we have a database with one binary relation $r \subseteq Name \times \{\top, \perp\} \times \{\top, \perp, *\}$, where *Name* is a set of names of some objects. The second column indicates whether an object is a bird and the third column indicates whether it flies. For example, the contents of the database might be the following:

Name	Bird	Flies
swallow	\top	\top
Tweety	\top	*
Clyde	*	*
Leo	*	\perp

We now assume that predicate symbols B and F correspond to the last two columns of the relation. Consider the following query $\Theta(Ab)$:

$$\forall x((B(x) \wedge \neg Ab(x)) \supset F(x)),$$

where Ab which stands for *abnormal* is the intensional predicate whose minimal definition we would like to compute. According to Theorem 4.2, the coherence condition for this query would be \top and the explicit definition for Ab would be $Ab(x) \equiv (B(x) \wedge \neg F(x))$, using the same method as we used for complete databases. Note that the relation F occurs positively in the query $\Theta(Ab)$. The calculated value of Ab would be \emptyset independently of the choice of predicates we choose to vary.

Now consider the query $\Theta'(N)$:

$$\forall x((B(x) \wedge N(x)) \supset F(x)),$$

where N which stands for *normal* is the intensional predicate whose *maximal* definition we would like to compute. The coherence condition for this query would again be \top and the explicit definition for N would be $N(x) \equiv (\neg B(x) \vee F(x))$. The following table then summarizes the possible results of calculating N .

B	F	Calculated value of N
varied	varied	$\{(swallow, \top, \top), (Tweety, \top, *), (Clyde, *, *), (Leo, *, \perp)\}$
varied	fixed	$\{(swallow, \top, \top), (Leo, *, \perp)\}$
fixed	varied	$\{(swallow, \top, \top), (Tweety, \top, *)\}$
fixed	fixed	$\{(swallow, \top, \top)\}$

■

It is important to note that since only the fourth step of the method for querying databases with complete information described in Section 4 differs from the querying method for querying databases with incomplete information described here, that the complexity results for steps one to three still hold. In addition, it is easily observed that step four for the new method does not add any new complexity to the querying method, so the complexity results apply to both querying databases with complete and incomplete information.

7 Conclusions

We have introduced a new query language SHQL and query method which we claim to be highly expressive, efficient, and natural to use. A PTIME querying method has been provided which is based on the use of quantifier elimination techniques. In addition, we have shown that the querying method may be used for databases with both complete and incomplete information. We believe the declarative character of the query language has much to offer in comparison to logic programming approaches, but at the same time is limited to only the class of PTIME queries. We are currently investigating extensions to the language and their characterization. As stated previously, because the approach is based on an existing algorithm (DLS*), there are already cases where queries outside the class we have investigated can be *compiled* and computed. We would also like to apply these techniques to commercial relational databases and are currently working on compilation methods from SHQL into standard SQL and extended SQL.

References

- [1] Ackermann, W. (1935) *Untersuchungen über das Eliminationsproblem der mathematischen Logik*, Mathematische Annalen, 110, 390-413.
- [2] Chandra, A. K. (1988) *Theory of Database Queries*, Proc. 7th ACM Symp. on Principles of Database Systems, 1-9.
- [3] Doherty, P., Lukaszewicz, W., Szalas, A. (1994) *Computing Circumscription Revisited. A Reduction Algorithm*, Technical Report LiTH-IDA-R-94-42, Linköping University, 1994. Preliminary report in: Proc. 14th IJCAI, 1995, Montreal, Canada.⁴ Full version to appear in Journal of Automated Reasoning.
- [4] Doherty, P., Lukaszewicz, W., Szalas, A. (1995) *A Characterization Result for Circumscribed Semi-Horn Formulas*, to appear in Fundamenta Informaticae. Also in: Technical Report LiTH-IDA-R-95-20, Linköping University, 1995.
- [5] Doherty, P., Lukaszewicz, W., Szalas, A. (1996) *General Domain Circumscription and its First-Order Reduction*, In: D. Gabbay and H.J. Ohlbach (eds.) *Proceedings of FAPR'96 Conference*, LNAI 1085, 93-109, Springer-Verlag. Extended version in: Technical Report LiTH-IDA-R-96-01, Linköping University, 1996.
- [6] Gurevich, Y., Shelah, S. (1986) *Fixpoint Extensions of First-Order Logic*, Annals of Pure and Applied Logic, 32, North Holland, 265-280.
- [7] Immerman, N. (1986) *Relational Queries Computable in Polynomial Time*, Information and Control, **68**, 86-104.
- [8] Immerman, N. (1989) *Descriptive and Computational Complexity*, Proceedings of Symposia in Applied Mathematics, 38, 75-91.
- [9] Kanellakis, P.C. (1990) *Elements of Relational Database Theory*. In: J. Van Leeuwen (ed.) *Handbook of Theoretical Computer Science*, vol. B, Elsevier, 1073-1156.
- [10] Nonnengart, A., Szalas, A. (1995) *A Fixpoint Approach to Second-Order Quantifier Elimination with Applications to Correspondence Theory*, to appear in *Logic at Work. Essays Dedicated to the Memory of Helena Rasiowa*, E. Orłowska (ed.), Kluwer, 1996. Also in: Report of Max-Planck-Institut für Informatik, MPI-I-95-2-007, Saarbrücken, Germany.
- [11] Reiter, R. (1978) *On Closed World Databases*, in: *Logic and Databases*, H. Gallaire and J. Minker (eds.), Plenum Press, New York, 55-76.
- [12] Sazonov, V. (1980) *A Logical Approach to the Problem of "P=NP"*, LNCS 88, Springer-Verlag, 562-575.
- [13] Szalas, A. (1993) *On the Correspondence Between Modal and Classical Logic: an Automated Approach*, Journal of Logic and Computation, **3**, 605-620.
- [14] Vardi, M. Y. (1982) *The Complexity of Relational Query Languages*, 14th ACM Symp. on Theory of Computing, 137-146.

⁴URL: <http://www.ida.liu.se/labs/kplab/people/patdo/>.