



Proceedings of the
XIII Spanish Conference on Programming
and Computer Languages
(PROLE 2013)

Towards Bridging the Expressiveness Gap Between
Relational and Deductive Databases

Fernando Sáenz-Pérez

22 pages

Towards Bridging the Expressiveness Gap Between Relational and Deductive Databases

Fernando Sáenz-Pérez^{1*}

Grupo de programación declarativa (GPD),
Dept. Ingeniería del Software e Inteligencia Artificial,
Universidad Complutense de Madrid, Spain¹

Abstract: SQL technology has evolved during last years, and systems are becoming more powerful and scalable. However, there exist yet some expressiveness limitations that can be otherwise overcome with inputs from deductive databases. This paper focuses on both practical and theoretical expressiveness issues in current SQL implementations that are overcome in the Datalog Educational System (DES), a deductive system which also includes extended SQL queries with respect to the SQL standard and current DBMS's. Also, as external database access and interoperability are allowed in DES, results from the deductive field can be tested on current DBMS's. For instance: Less-limited SQL formulations as non-linear recursive queries, novel features as hypothetical queries, and other query languages as Datalog and Extended Relational Algebra. In addition, some notes on performance are taken.

Keywords: Relational Databases, Deductive Databases, SQL, Datalog, Expressiveness

1 Introduction

Deductive database systems extend (“relational”) database management systems (DBMS's) by including a more powerful query language based on logic. Datalog (and its extensions as Datalog with negation, uninterpreted function symbols, disjunctive heads, constraints, ... [RU93]) became the *de-facto* deductive query language. This language has been extensively studied and nowadays is regaining an increased interest in different database application areas [Got12].

Motivations for researching in deductive databases include clean semantics of logic-based approaches, neat formulations, and both expressiveness and performance gains [TLLP08]. Provided that the deductive data model (function-free case) meets the relational one with respect to the data structures, it is possible to query the same database (either intensional or extensional, and either deductive or relational) with different languages, as Datalog, SQL, and Relational Algebra (RA).

Language expressiveness can be seen from two points of view: Expressive power (*theoretical* expressiveness), and concisely and readily (*practical* expressiveness). Nowadays, one can face several obstacles when formulating SQL queries because of a number of reasons, including, e.g.:

* This author has been partially supported by the Spanish projects CAVI-ART (TIN2013-44742-C4-3-R), STAMP (TIN2008-06622-C03-01), Prometidos-CM (S2009TIC-1465) and GPD (UCM-BSCH-GR35/10-A-910502)

parsing requirements and restricted syntax constructions for practical expressiveness, and limited set of operators and recursion limitations for theoretical expressiveness. Whereas the user can work around practical expressiveness obstacles, usually contrived formulations are reached. On the other hand, theoretical expressiveness even makes impossible to formulate a given query.

In this work, an ongoing work is presented, enabling more powerful and less restrictive forms of SQL queries (w.r.t. current DBMS's) in the deductive system DES (Datalog Educational System) [SCG11], a system targeted at teaching SQL, RA and Datalog at class rooms. Nonetheless, it is also used for many purposes (cf. DES Facts at [SP13a]). It can be argued that less restrictive SQL queries enable students a more rapid learning curve because less problems arise in writing such queries. Also, we consider extended SQL, RA and Datalog languages with novel features which have not been presented before as functional dependencies, the division (relational algebra) operator, SQL hypothetical queries, and Datalog hypothetical queries and rules. Along the paper, for some SQL queries, there are presented equivalent Datalog and RA queries which highlight equivalent formulations in these languages. In this paper we also present for the first time persistence and database interoperability as supported by DES, which allows to test new language features (less-limited SQL recursion, hypothetical queries, division operator, ...) in current relational database systems. Although this is not the first deductive database implementing persistence (cf. [TLLP08, AOT⁺03]), it develops a seamless integration in a truly interactive system with external DBMS's.

Organization of this paper proceeds as follows. In Section 2, DES is briefly introduced. Next, Section 3 describes the various scenarios for accessing external DBMS's and mixed query solving in DES. Section 4 focuses mainly on the limitations of state-of-the-art SQL-based DBMS's, which can be overcome with a deductive database, also illustrating equivalent formulations in Datalog and RA, and introducing hypothetical queries as an addendum from deductive databases. Although DES is not geared towards performance, Section 5 briefly analyzes some benchmarks that illustrate that a competitive system might be achieved. Finally, Section 6 concludes and points out some future work.

2 Datalog Educational System

DES is a free, open-source, interactive, multiplatform, portable, Prolog-based implementation of a deductive database system. DES 3.4 [SP13a] is the current implementation, which enjoys Datalog, (extended) Relational Algebra (RA) and SQL query languages, persistence, full recursive evaluation with tabling, full-fledged arithmetic, strong constraints, stratified negation as described in [UII88] with safety checks [UII88, ZCF⁺97], ODBC connections, and novel approaches to Datalog and SQL declarative debugging [CGS12, CGS11, CGS08], test case generation for SQL views [CGS10], duplicates, null values and outer join support [SP12], aggregate predicates and functions [SP12], and hypothetical queries and rules. It is a live system experiencing many downloads (>49K) and used all over the world both for teaching and research (cf. DES Facts at [SP13a]). Revealing itself not only interesting to learn deductive databases, feedback from teachers suggest that even logic programming is more well-understood by the students with the more declarative language Datalog than Prolog (cf. Quotes at [SP13a]). Indeed, pure Datalog does not include non-declarative constructs as Prolog does. For instance, the order of

clauses and goals in Prolog can affect to finding answers because of non-terminating branches. Also, the cut operator can discard actual answers. However, such issues are not the case for pure Datalog.

Interacting with DES is possible via either an OS command shell or GUI applications, as the Java-based IDE ACIDE [SP07], TextMate, Emacs, Crimson Editor and others. As the system is implemented on top of Prolog, it can be run from a state-of-the-art Prolog interpreter (currently, last versions of SWI-Prolog and SICStus Prolog are supported) on any OS supported by such Prolog interpreter. Portable executables (i.e., they do not need installation and can be run from any directory they are stored) have been also provided for Windows, Linux, and Mac OS X.

There is available a wide set of commands for dealing with the system (in-memory database, ODBC connections, debugging and test case generation, OS interaction, persistence, etc.) They are preceded by a slash to syntactically isolate commands and queries. Assertions are also provided, as, e.g., type constraints, so that typed Datalog predicates are available as relations for SQL to be queried (SQL require typed relations, both tables typed by the user and views whose types are automatically inferred).

Datalog as supported by DES mainly follows Prolog ISO standard [ISO00] (considering its syntax as a subset of Prolog), whilst SQL follows SQL:2008 ISO standard [ISO08], and RA follows [Cod72] extended with recursion, nulls, outer joins and aggregates (syntax is borrowed from [Die01]). Their concrete syntax can be found at [SP13a]. All of these query languages can access the very same database, which is implemented with both an in-memory Prolog database, and external relational databases (cf. Section 3). SQL statements that are solved by DES are compiled to Datalog rules, as well as all RA statements. The command `/show_compilations on` enables the display of such compilations.

Evaluation of Datalog recursive queries is ensured to be terminating if two conditions are met. First, no infinite predicates/operators are considered. Currently, only the infix operator “is” (and its more general form of equality “=”) represents an infinite relation and can deliver unlimited pairs (other built-ins, as comparison operators, demand their arguments to be ground). For instance, let’s consider the rules $p(0)$ and $p(X) :- p(Y), X \text{ is } Y+1$. Then, the query $p(X)$ is not terminating since its meaning is infinite ($\{p(0), p(1), \dots\}$). Note, however, that a Datalog novel top-N query can be submitted, analogously to some current DBMS’s, as $\text{top}(10, p(X))$, which finds for the first ten answers and does terminate. Second, duplicates are not enabled for hypothetical queries [SP13b, SP13a]. For example, let’s consider solving the query p in the context of the (single-rule) program $p :- p \Rightarrow p$. With duplicates enabled, this rule means that p is true if, assuming a *new* fact p (to the left of \Rightarrow) then p (to the right of \Rightarrow) must be true. Each time this rule is selected for proving p , the assumed fact for the database is considered as a new instance, different from others either already in the EDB or previously assumed. Therefore, as no terminating condition is located at the rule, non-termination occurs. Recalling top-N queries, they are ensured to terminate if two conditions hold: First, the very same condition for hypothetical queries, and, second, limiting the relations involved in the proof of the query to belong to the same stratum (cf. [UII88]). This second condition avoids the situation for which the complete meaning of a (possibly infinite) relation in a deeper stratum is needed. Negation, aggregates, and ordering (meta)predicates require to set their relational arguments to be in a deeper stratum than the rule in which they occur. As in other systems, such as DLV and relational databases including recursion (e.g., DB2), including parametric bounds either for the

domains of numbers or fixpoint iterations can be considered to limit non-terminating queries as a subject of future work.

Datalog temporary views allow to write compound queries on the fly (as, e.g., conjunctions and disjunctions). A temporary view is a rule which is added to the database, and its head is submitted as a query and executed. Afterwards, the rule is removed. For instance, given the relations $a/1$ and $b/1$, the temporary view $d(X) :- a(X), \text{not}(b(X))$ computes the set difference between the instance relations a and b .

As it will be explained in the next section, SQL statements can be either translated into Datalog programs and executed by the deductive engine, or sent to and solved by an external DBMS. Also, SQL statements can be the result of translating Datalog persistent predicates in order to be externally solved by a DBMS.

3 Enabling Interoperability

Figure 1 depicts the system architecture of DES supporting this functionality. Datalog queries are solved by the deductive engine relying on a cache to store results from fixpoint computations by using a tabling technique [Die]. This technique allows to detect already computed calls and avoid to recompute them by reusing their stored computed answers. The deductive engine is able to solve such queries and to pass SQL queries (as input by the user or as a result of predicate persistence) to external DBMS's. Next, the different scenarios for in-memory, external, and mixed query solving are presented.

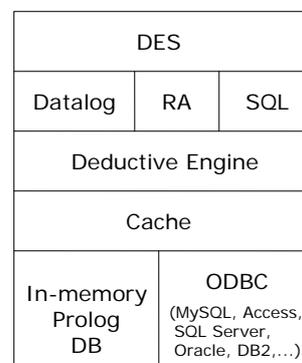


Figure 1: System Architecture

3.1 In-Memory Database

First alternative for dealing with SQL statements is to use the very same Prolog database. SQL row-returning queries for the in-memory database are translated into and executed as Datalog programs by the deductive engine (basics can be found in [UII88]), and relational metadata for DDL (Data Definition Language) statements are kept. Submitting a DML (Data Manipulation Language) row-returning query amounts to 1) parse it, 2) translate into a Datalog program including the relation $answer/n$ with as many arguments as expected from the SQL statement and including the translated Datalog form corresponding to the SQL statement, 3) assert this program, and 4) submit the Datalog query $answer(X_1, \dots, X_n)$, where $X_i : i \in \{1, \dots, n\}$ are n fresh variables. After its execution, this translated Datalog program is removed. On the contrary, if a DDL statement defining a view is submitted, its translated Datalog program and metadata do persist. This allows Datalog programs to seamlessly use tables and views created at the SQL side. Note that SQL metadata (column names and types) are stored as Datalog assertions, and all consistency constraints (as primary keys and referential integrity constraints) are also supported as Datalog assertions.

As an example of a DDL statement, let's consider:

```
create table employee(name varchar, dept varchar, salary integer);
```

which is equivalent (and can be interchangeably used) to the following Datalog assertion (which is denoted by a preceding neck symbol “:-”):

```
:-type(employee(name:varchar, dept:varchar, salary:integer)).
```

Also, integrity constraints can be stated, as the following primary key assertion (or added as a either a column or table constraint in the former `create table` statement):

```
:-pk(employee, [name]).
```

As the result of submitting the former inputs, we can inspect metadata for the current database (here, the in-memory, default database `$des`) with:

```
DES> /db_schema
Info: Database '$des'
Info: Table(s):
  * employee(name:string(varchar),
             dept:string(varchar),salary:number(integer))
    - PK: [name]
Info: No views.
Info: No integrity constraints.
```

Inserting tuples can be done with either an SQL `insert` statement or a command:

```
insert into employee values('Smith','Sales',15000);
/assert employee('Smith','Sales',15000).
```

Selecting tuples can be performed with either an SQL `select`, or a Datalog query, or an RA query. In the following example, employees with a salary over 10,000 are requested for these query languages, respectively:

```
DES> select * from employee where salary > 10000;
answer(employee.name:string(varchar),employee.dept:
string(varchar),employee.salary:number(integer)) ->
{
  answer('Smith','Sales',15000)
}

DES> employee(N,D,S), S>10000
{
  employee('Smith','Sales',15000)
}

DES> select salary>10000 (employee)
answer(employee.name:string(varchar),employee.dept:
string(varchar),employee.salary:number(integer)) ->
{
  answer('Smith','Sales',15000)
}
```

As DES Datalog implementation follows Prolog syntax, variable names start with uppercase (as `N`, `D`, and `S` in the query) and user identifiers start with lowercase (as `employee`). This also applies to RA variable and user identifier names.

Although this in-memory approach is not understood as a persistent database, it is still possible to save the contents of the current database to a file and afterwards recover them with the commands `/save_ddb`, and `/restore_ddb`, respectively. Next two sections deal with other approaches relying on external DBMS's for dealing with persistent data.

3.2 Connecting to External Databases

An ODBC connection is identified by a name defined at the OS level, and opening a connection in DES means to make it the current database. Any relation defined in the external DBMS as a view or table is allowed as any other relation (predicate) in the deductive database. So, this second alternative to dealing with SQL statements is to use the ODBC bridge to access external databases: Each SQL relation (table or view) is understood as a predicate and therefore can be seamlessly accessed from either a Datalog or an RA query. Contents of SQL relations are retrieved (possibly involving query processing in the case of views) from the current open external DBMS, and SQL queries are directly sent to and processed by such DBMS. As both Datalog and RA queries can refer to SQL relations, therefore, computing such a query can involve computations both in the deductive inference engine and in the external SQL engine. For instance, let's consider the external table `manager(mgr varchar, emp varchar)` stating that `mgr` is the direct manager of `emp`. The following Datalog query computes all managers (both direct and indirect):

```
DES> /open_db mysql
DES> create table manager(mgr varchar(10), emp varchar(10));
DES> insert into manager values ('E1','E2'), ('E2','E3'), ('E2','E4');
DES> managers(M,E) :- manager(M,E) ;
                        manager(M,M2), managers(M2,E) .

Info: Processing:
      managers(M,E)
in the program context of the exploded query:
      managers(M,E) :- manager(M,E) .
      managers(M,E) :- manager(M,M2), managers(M2,E) .
{
  managers('E1','E2'), managers('E1','E3'), managers('E1','E4'),
  managers('E2','E3'), managers('E2','E4')
}
```

This is a Datalog temporary view which recursively computes the outcome as the union (expressed with the semicolon) of direct managers (`manager(M,E)`) and managers which has also other managers (`manager(M,M2), managers(M2,E)`). Note also that the system informs about how this is translated into a query (`managers(M,E)`) and a set of rules (exploded query). Whereas the Datalog engine is responsible of computing the transitive closure, the external SQL engine provides the data source for the relation `manager`. Note that in this example the external computation is rather light, but consider that this relation can be a complex view or we are looking for a concrete employee's managers. In this last case, e.g., the submitted SQL query to the external engine would include the `where` condition.

Using SQL is also possible for answering the same transitive closure and referring to an external database, even when such database does not support recursion. Recall, however, that when the current database is an ODBC connection, SQL statements are directly injected to the external database. To allow DES to process such query with its own engine, the command `/des_sql_solving on` is provided. This command instruct DES to override the default behavior and otherwise translate SQL statements into Datalog rules, and solve these with its engine, in a similar way as the managers example proceeds. In addition, the schema of `$des` is also made visible, as temporary views can be defined. Next example shows all this when the current

database is `mysql` (the result of the query is the same as before but with the default predicate name for the answer).

```
DES> /des_sql_solving on
DES> with managers(mgr,emp) as
      select * from manager union
      select manager.mgr,managers.emp from manager, managers
      where manager.emp=managers.mgr
      select * from managers;
Info: SQL statement compiled to:

managers(A,B) :-
  distinct(managers_2_1(A,B)).
managers_2_1(A,B) :-
  manager(A,B).
managers_2_1(A,B) :-
  manager(A,C),
  managers(C,B).

Info: SQL statement compiled to:

answer(A,B) :-
  managers(A,B).

answer(managers.mgr:string(varchar(10)),managers.emp:string(varchar(10))) ->
{
  answer('E1','E2'), answer('E1','E3'), answer('E1','E4'),
  answer('E2','E3'), answer('E2','E4')
}
```

Submitting RA statements is also possible in this scenario. The following system session illustrates the same example above forcing DES to evaluate a query as an RA query (with the command `/ra`; otherwise it is interpreted as an SQL query and it is syntactically rejected), and also listing the compilation result (command `/show_compilations on`)¹:

```
DES> /show_compilations on
DES> managers(emp,mgr) :=
      select true (manager) union
      project manager.mgr,managers.emp
      (select manager.emp=managers.mgr
       (manager product managers));
Info: RA expression compiled to:
managers(A,B) :-
  manager(A,B).
managers(A,B) :-
  manager(A,C),
  managers(B,C).
DES> /ra select true (managers)
```

Concluding, opening an ODBC connection allows the user to integrate external relations in Datalog, (DES-solved) SQL, and RA queries, but in this setting it is not possible to integrate Datalog predicates in external SQL queries, as the external database is not aware of Datalog relations. To make it possible, another alternative is presented next.

¹ The result is omitted as it is the same as the previous example.

3.3 Persisting Predicates

Persisting a predicate in DES allows durability for deductive programs by relying on current relational DBMS's as persistent media. We have proposed an assertion as a basic declaration for making a predicate to persist, similar to [CGC⁺04]. Such work implements extensional predicate persistence in a Prolog system, i.e., only predicates composed of facts can be persisted. In contrast, DES also allow persistent intensional predicates (including both facts and rules). The general form of a persistence assertion in DES at the command prompt is as follows:

```
:- persistent (PredSpec[, Connection]).
```

where `persistent` is the keyword for enabling persistence, `PredSpec` is a predicate schema specification, and the optional argument `Connection` is an ODBC connection identifier. Such schema specification can be either `PredName/Arity` or `PredName(Schema)`, where `Schema` can be either `ArgName1, ..., ArgNamen` or `ArgName1:Type1, ..., ArgNamen:Typen`. If a connection name is not provided, the current open database is used (the local, default database `$des` cannot be used for persistence). With this assertion, we allow for: First, persisting both an empty predicate (i.e., with no defining rules) and an already defined predicates (with defining rules). And, second, for persisting both an untyped or already typed predicate.

The following example makes persistent the predicate `employee`, as already defined in the deductive database in the previous subsection. The second argument (`mysql`) is the name of the ODBC connection (in this case, MySQL is the target DBMS) to map the predicate.

```
:- persistent (employee/3,mysql).
```

Any rule belonging to the definition of a predicate `p` which is being made persistent is expected, in general, to involve calls to other predicates. Each callee (such other called predicate) can be:

- An existing relation in the external database.
- An already persisted predicate which is loaded in the local database.
- An already persisted predicate which is not yet loaded in the local database.
- A predicate which has not been made persistent yet.

For the first two cases, besides making `p` persistent, nothing else is performed when processing its persistence assertion. For the third case, a persistent predicate is automatically restored in the local database, i.e., it is made available to the deductive engine. For the fourth case, each non-persistent predicate is automatically made persistent if types match; otherwise, an error is raised. This is needed in order for the external database to be aware of a predicate which is only known by the deductive engine so far, as this database will eventually compute the meaning of `p`.

However, not all persistent rules are processed by the external DBMS because it does not support some features, and the translations of some built-ins are not supported yet. In the current state of the implementation, the following conditions must hold for a rule to be externally processed:

- The rule does not contain calls to built-ins but comparison operators.

- The rule is not included in a recursive cycle.

Rules that do not meet these conditions are denoted as *non-projected* rules. Anyway, non-projected rules are externally stored as metadata information for their persistence, so that it is possible to restore the complete predicate in different sessions. Non-projected rules are also kept in the in-memory database for computing the meaning of the predicate when requested. This is performed by the deductive engine, which couples the processing of the external database with its own processing to derive the meaning of the predicate. Therefore, all the deductive computing power is preserved although the external persistent media lacks some features as, for instance, recursion (see next section for an example). Further releases might contain relaxed conditions for this translation stage as for, e.g., allowing to project those recursive SQL queries that are supported by the external DBMS. Note that persistent predicates can be used as a regular SQL relation in statements directly submitted in the external DBMS. However, only for the predicates with all the rules meeting the conditions above, it is possible to compute the same set of tuples as DES does because the non-projected rules are not known to the external SQL engine.

3.4 Intermixing Query Solving

As introduced before and detailed in [SP12], DES uses an inference engine for solving Datalog queries. To this end, a fixpoint computation by strata is conducted to deliver the answer to a given query, keeping already-made calls and already-computed answers as tables in memory. (To enhance performance, we use hash indexing for these different tables.) Along fixpoint iterations, program rules are selected in a top-down fashion for solving the query, i.e., looking only for the needed rules relevant to the query. In particular, those relations that are not in the sub-PDG (Predicate Dependency Graph [UII88]) for the query does not participate in its solving, and moreover only matching rules (i.e., those whose head that unify with a given call during solving) for the relevant relations are used.

Given this, the backbone for intermixing query solving lies on the data provider. When considering a (non-persistent) Datalog relation, each one of its defining rules (both extensional and intensional) are managed by the underlying Prolog (in-memory) dynamic database. So, requesting for a matching rule for a given call during solving involves to ask the underlying Prolog system to find it in its database. Depending on the Prolog host, this usually involves to sequentially scan the in-memory database for matching rules unless implicit indexing is available (as in SWI-Prolog). Representing a Datalog rule with a Prolog predicate is via `datalog/7`, a predicate having the Datalog rule as its first argument, and containing several other arguments for holding extra information (such as variable names, rule identifier, hypothetical context [SP13b], rule source, ...). The following is an instance of this predicate holding two facts (`t(1)` and `s(2)`), and a rule (`t(X) :- s(X)`):

```
datalog(t(1), [], 0, [], [], asserted((2014,1,9,15,15,55)), source).
datalog(s(2), [], 1, [], [], asserted((2014,1,9,15,16,1)), source).
datalog((t(A) :- s(A)), ['X'=A], 2, [], [], asserted((2014,1,9,15,18,23)), source).
```

When an ODBC connection is the current database, on the one hand, any SQL query is directly injected to the external database, so that query solving completely relies on the external relational DBMS (recall that in this mode, no references to Datalog relations are possible from an SQL

query). A set of glue predicates are implemented to collect data and schema from the external DBMS and to be processed for the display of the query answer at the DES system prompt. On the other hand, if the submitted query is a Datalog query, it is solved as before, using the rules stored in the Prolog dynamic database. However, as opposed to SQL queries with an ODBC connection as the current database, this query can contain references to SQL relations stored in the external DBMS, which will be in charge of computing such relations. So, to allow intermixed query solving involving both the deductive engine and the external SQL engine, the Datalog rule database is overloaded with a clause that requests data from the external DBMS:

```
datalog(Rule, [], RuleId, [], [], rdb(Connection), source) :-
    datalog_rdb(Rule, [], RuleId, [], rdb(Connection), source).
```

Here, `rdb(Connection)` refers to the external relational database for which a given ODBC connection name is given at the OS level. When during query solving a given call is performed, the Datalog rule base is scanned, but not only obtaining the rules in the dynamic database (Prolog `datalog facts`) as before, but also by using the external data collected with the added clause just shown. This clause calls the predicate `datalog_rdb`, which builds an SQL statement for retrieving matching data from the external DBMS. As matching rules are selected by unifying the call with the head of the rule, it is possible to have arguments bound to constants (including a null value). Therefore, a `WHERE` SQL condition is added to the statement filtering data as the ground arguments require. For instance, if the call is `p(1, 'str', X)` (and assuming a table `p` with schema `(a integer, b string, c string)`), then the generated SQL is:

```
select a from p where a=1 and b='str';
```

Note that no condition is added for the third argument, as it is not ground. If an argument, say `a`, comes bound to `null`, then the condition turns to `a is null`, as the SQL equality operator cannot be used to test for null values. Such external data providers are obviously only tried for table and view names that exist already in the external database and matching the predicate call (`t` in the example).

Retrieving data from both sources means that the meaning of an external relation is computed as the set (or multiset, if duplicates are enabled) union of the Datalog relation and the external relation. Assuming that the external DBMS with ODBC identifier `db2` contains a table `t` with tuples `{ (3), (4) }`, and considering a Datalog predicate with tuples `{ (1), (2) }`, then submitting the following queries give the following results (where the command `/use_db` sets the current database to its argument):

```
DES> /use_db $des
DES> t(X)
{
  t(1), t(2)
}
DES> /use_db db2
DES> t(X)
{
  t(1), t(2), t(3), t(4)
}
DES> select * from t
answer(a:INTEGER(4)) ->
```

```

{
  answer(3), answer(4)
}
DES> /des select * from t
Info: SQL statement compiled to:
answer(A) :-
  t(A).
answer(t.A:number(integer)) ->
{
  answer(1), answer(2), answer(3), answer(4)
}

```

So, querying to the default deductive database only retrieves its data, but when opening the external database it retrieves data from both sources. The SQL query directly injected to the external database retrieves data only from the external database, but forcing DES to solve the same query (using the alternative command `/des` for a single query) data from both sources are also retrieved.

Finally, we deal with the case of persistent predicates. To deal with such predicates, another clause for the predicate `datalog/7` is added, and it proceeds similar to the handling of external DBMS relations, because for a persistent predicate, a corresponding view (with the same name as the predicate) is created in the DBMS. For a predicate `p`, this view is created as the union of a table `p_des_table` (storing the predicate extensional rules, i.e., its facts) and the equivalent SQL query for the remaining intensional rules (i.e., its rules with both head and body).

```
CREATE VIEW p AS SELECT * FROM p_des_table UNION DL_to_SQL(rules(p))
```

where `rules(p)` are the intensional rules of `p`, and `DL_to_SQL` is the function that translates a set of Datalog rules into an SQL query, providing the same meaning than the set of rules. To implement this function, we have resorted to Draxler's Prolog to SQL compiler [DS92], and the function is replaced by its result in the SQL sentence sent to the DBMS for the creation of the view.

A persistent predicate can be queried either from the default deductive database (`$des`) or from the ODBC database where it has been made persistent. However, and as already introduced, not all persistent Datalog rules can be processed by the external DBMS because either it does not support some features (e.g., hypothetical reasoning and non-linear recursion) or the translations of some built-ins are not supported yet. In such a case, the deductive engine couples its own processing with the processing of the external database in the following way. Let a predicate `p` be defined by a set of rules `S` that can be externally processed and other set of rules `D` that cannot. Then, the meaning of `p` ($\| p \|$) is computed as the union of the meanings of both sets of rules computed by the corresponding engines (external SQL - *SQL* - and Datalog - *DL*):

$$\| p \| = \| DL_to_SQL(S) \|_{SQL} \cup \| D \|_{DL}$$

following an analogous process to the one depicted above, where the call `t(X)` included data from both the external DBMS and the Datalog dynamic database. Note that this allows to retain all the expressive power of Datalog in defining relations, so that it is possible to use features absent in the external database as, e.g., hypothetical reasoning (see next section).

4 Extending DBMS Expressiveness

This section shows some limits of both expressiveness types (both practical and theoretical) in current DBMS's and which are otherwise overcome with DES. All (SQL, Datalog and RA) queries presented along this section have been actually computed in this system. This, coupled with interoperability, can allow users to use a less-restricted form of SQL and test new features with their DBMS of choice.

4.1 Practical Expressiveness

Here, some sources of practical expressiveness limitations in current DBMS's are highlighted, which are all of them overcome by DES. First, some parsers require to write extended formulations even when it would not be strictly needed. For instance, selecting all arguments from an autojoin is not possible in Access and MySQL as:

```
select * from t,t;
```

and table renamings must be added (MySQL requires each derived table to have its own alias).

Some DBMS's as Sybase, MySQL and Access do not include all the outer join versions. Let's consider a full outer join:

```
select * from s full join q on s.sno=q.sno;
```

Therefore, to compute a full outer join one must resort to add more code, as, e.g.:

```
select * from s left join q on s.sno=q.sno
union all
select * from s right join q on s.sno=q.sno;
```

If relations s and q are queries instead of relations, this becomes worse as one has to either duplicate code for them or create new views (with associated issues as changing the database, administration privileges, and hiding original queries).

Also, many nested uses of outer joins are rejected for a number of reasons (including issues related to univocally identify relations and columns in correlations) in different DBMS's (Access, DB2, MySQL, Oracle, PostgreSQL, SQL Server and Sybase) such as:

```
select * from s left join
(select * from q right join sp on q.sno=sp.sno where q.qname<sp.pname)
on s.sno=q.sno where s.name=q.qname;
```

Finally, the absence of standard operators as `EXCEPT` (e.g., MySQL and Access) makes the formulation of several queries more cumbersome and less efficient. Let's assume the query:

```
(select * from s) EXCEPT (select * from t);
```

Without this operator, the query can be expressed with `NOT IN`, but requires to explicitly include the arguments in the schema of the left relation, as in:

```
select * from s where (s.s1,...,s.sn) not in (select * from t);
```

Even worse, not all DBMS's allow to use tuples as elements in the containment operation, so that the more cumbersome `NOT EXISTS` with correlations should be used:

```
select * from s where (s.s1,...,s.sn) not exists
(select * from t where s.s1=t.t1 and ... and s.sn=t.tn);
```

However, DES does support it and allow all of the above neater formulations.

4.2 Recursive SQL

Here, we consider some issues regarding theoretical expressiveness for recursive queries. Let's consider a classical transitive closure problem: Given a graph defined by `edge(ori, des)`, find all paths. A possible, simple, recursive SQL formulation follows:

```
with recursive path(ori,des) as
  select edge.*,1 from edge
union
  select path.ori,edge.des from path,edge where path.des=edge.ori
select * from path;
```

But it is rejected in DB2, Oracle, PostgreSQL and SQL Server because they disallow the recursive keyword and/or require `union all`, and parentheses around the local view definition (in turn, neither Access nor MySQL supports SQL recursion). In fact, `union all` is a requisite in standard SQL because of discarding duplicates. However, DES does not require this and both a set or multiset answer can be requested. Also, thanks to the support of Datalog, the following neater formulation is allowed:

```
path(ori,des) :-
  edge(ori,des)
;
edge(ori,Int), path(Int,des).
```

Note that, as the former SQL `with` formulation, this Datalog rule is also a query (a temporary view) which can be submitted at the command prompt. In contrast to current DBMS's, this SQL query can be included in a simplified view declaration in DES as follows, avoiding the use of the `with` clause and additional relation names:

```
create view path(ori,des) as
  select edge.* from edge
union
  select path.ori,edge.des from path,edge where path.des=edge.ori;
```

Another problem with this SQL query in current DBMS's is when the graph includes a cycle. Oracle simply rejects such queries even when they can be actually computed, PostgreSQL enters an infinite loop, and DB2 shows top-500 results by default (requiring all tuples also yields to non-termination). Note that DES, even when enabling duplicates (duplicate sources are both extensional and intensional definitions [SP12]), is able to terminate in this situation because, in contrast to those systems, checks already solved calls along fixpoint iterations, avoiding to recompute them. This feature is possible thanks to the tabling technique, which holds both already-computed answers and already-tried calls for avoiding to recompute subsumed calls. However, if a query looking for paths involves also their lengths, this technique will not help as the answer would be infinite in the presence of cycles in the graph (because of the infinite relations due to arithmetics). For instance, with `edge={ (a,b), (b,a) }`, paths can be constructed visiting `a`, then `b`, returning to `a`, and so forth, with different lengths $((a,b,1), (a,a,2), (a,b,3), \dots)$

Current DBMS's show another limitation: Linear recursion is required, which means that only one recursive call is allowed in a recursive definition. The following system session shows, first, how relations `edge` and `path` are made persistent and, second, that non-linear recursion

is allowed in DES. To this end, processing of the recursive, non-linear rule is conducted by the deductive engine.

```
:-persistent (edge(a:int,b:int),mysql).
:-persistent (path(a:int,b:int),mysql).
with recursive path(a, b) as
  select * from edge
  union
  select p1.a,p2.b from path p1, path p2 where p1.b=p2.a
select * from path;
```

This example follows previous ones and, as seen, can be formulated as linear-recursive. However, without non-linearity, several queries cannot be expressed, as for instance some graph algorithms [ZCF⁺97]. Other examples, as Fibonacci, requires to find a reformulation for its mathematical specification, therefore revealing a practical expressiveness limitation. Also, recursive SQL queries involving `EXCEPT`, `NOT IN`, and aggregates are not allowed in current DBMS as termination is neither ensured nor detected. But deductive systems implementing XY-stratification can solve some of them, as LDL++ [AOT⁺03] (current DES system does not implement this yet).

Mutual recursion is also appealing to formulate some relations, as the hubs and authorities example or the following classical definition for even and odd relations:

```
with
  even(x) as select 0 union select odd.x+1 from odd,
  odd(x) as select even.x+1 from even
select top 20 x from odd;
```

Note that `select 0` is a from-less SQL statement as accepted in several DBMS's, which simply returns a tuple as specified in the projection list (in this case: (0)). However, mutual recursion is not allowed in current DBMS's but otherwise allowed in DES.

4.3 The Division RA Operator

An overlooked, but important, relational algebra operator is division, as found in the original proposal of Codd [Cod72]. Although attracting interest [con06], no current DBMS includes this operator, which identifies the attribute values from a relation that match with all of the values from another relation. There are several approaches to solve this kind of queries [McC03], but they resort to cumbersome formulations that can be avoided if this operator was made available. DES includes a novel syntax for allowing the division operator in SQL by extending the form that a relation may take as follows (in BNF):

```
Relation ::= ... | Relation DIVISION Relation
```

Consider the subset of Date's [Dat09] famous SuppliersPartsProjects schema $p(pno, pname, color, weight, city)$ and $spj(sno, pno, jno, qty)$, for parts and suppliers providing parts, where sno and pno stand for supplier and part identifiers, respectively. We are interested in the query [McC03]: "Find the sno values of the suppliers that supply all parts of weight equal to 17." This can be easily formulated in DES as:

```
select * from
  select sno,pno from spj
  division
  select pno from p where weight=17;
```

One can compare this neater formulation to those contrived, lengthly and error-prone equivalent formulations in [McC03]) as, for instance, the direct translation from its definition into relational algebra (assuming that the EXCEPT operator is available):

```
select distinct sno from spj
except
select sno
  from (select sno, pno
        from (select sno from spj) as t1,
             (select pno from p where weight=17) as t2
        except
         select sno, pno from spj
       ) as t3;
```

The division operator is also available in DES in both RA and Datalog as a novelty, and the very same query can be respectively issued in both languages as:

```
(project sno,pno (spj)) division (project pno (select weight=17 (p)))
spj(SNO,PNO,_,_,_) division p(PNO,_,_,17,_)
```

A couple of notes in this example are: First, the division operator in Datalog refers to variable names, instead of column names in a schema. And, second, whereas in RA and SQL, projection is needed to build the appropriate schema of involved operands of the division operator, in Datalog it suffices to denote non-relevant variables as anonymous. This allows to neatly discard those non-relevant variables from the schema of the operands for the division operator.

Datalog, SQL and RA statements including the division operator are translated into a set of Datalog rules implementing the direct RA formulation.

4.4 Functional Dependencies

Functional dependencies are key concepts in learning normalization, but neither standard SQL nor any current DBMS provide a way to enforcing them. Only DB2 supplies the construction DETERMINED BY as either a table or column constraint for data cubes. However, such functional dependency is only used to produce an optimized query plan, but it is not possible to enforce it yet. DES follows the syntax in DB2 allowing its enforcing with the following syntax, where `type_ctr` is a type constraint, and `colfi` and `colti` are column names:

```
colf1 type_ctr determined by colt2 -- column constraint
(colf1,...,colfn1) determined by (colt1,...,coltm) -- table constraint
```

which is equivalent to the also supported Datalog assertion syntax:

```
:-fd(Relation, [colf1, ..., colfn], [colt1, ..., coltm])
```

For example, given the table `offices(branch, street, city, zip)` stating the location of branches in different cities, the functional dependency $\{zip\} \rightarrow \{city\}$ (i.e., all the tuples with a given a zip postal code must have the same city value) can be declared and tested as follows:

```
DES> create table offices(branch string primary key,
    city string determined by zip, zip string);
DES> insert into offices values ('Major','New York','10007');
DES> insert into offices values ('Subsidiary','New York','10007');
DES> insert into offices values ('Admin','Washington D.C.','10007');
Error: Functional dependency violation offices.[zip]->offices.[city]
    in table offices(branch,city,zip)
    when trying to insert: offices(Admin,Washington D.C.,10007)
    Witness tuple          : offices(Major,New York,10007)
```

In contrast to current DBMS implementations, DES allow to add constraints on views (in particular, functional dependencies). Following the same example:

```
DES> create view v as select * from offices;
DES> alter table offices drop constraint check city determined by zip;
DES> alter table v add constraint check city determined by zip;
DES> insert into offices values ('Admin','Washington D.C.','10007')
DES> /check_db
Error: Functional dependency violation v.[zip]->v.[city]
    Offending values in database: [fd(Admin,Washington D.C.,10007),
    fd(Major,New York,10007),fd(Subsidiary,New York,10007)]
Error: In constraint: :-fd(v,[zip],[city])
```

Views are checked for consistency on demand, which is possible with the command `/check_db`.

4.5 Hypothetical SQL Queries

As a novel feature, DES includes hypothetical SQL queries (absent in the standard) for solving “what-if” scenarios. Date [Dat09] explains this idea (firstly proposed in [SK80]) which broadly means that data can be assumed for a given query without actually modifying its source relations. DES includes a novel syntax for allowing such assumptions in the form of:

```
assume SQL_stmt1 in Rel1 [, ...,SQL_stmtN in RelN]
SQL_query;
```

which allows to assume the result of `SQL_stmti` in `Reli` when processing `SQL_query`. This syntax for hypothetical SQL clauses follows a similar syntax of `with` clauses, but using `assume ... in` instead of `with ... as`. Indeed, while a `with` clause allows to *define* new temporary relations with queries, an `assume` clause allows to *modify* the semantics of already-defined relations.

Referring to the former example about recursive paths, one might be interested in assuming that a given edge (say `(3,1)`) is in the relation `edge` in order to know the new transitive closure of `path` without resorting to update `edge`. Following the syntax above, this can be formulated as follows, which assumes another (extensional) edge in the `path` recursive view as defined before:

```
assume select 3,1 in edge select * from path;
```

Also, intensional, recursive rules can be assumed, as in the following query, which computes the transitive closure of `edge` (the relation `path` is not used):

```
assume
  select e1.ori, e2.des from edge e1, edge e2 where e1.des=e2.ori
in edge
select * from edge;
```

Although the current version of DES restricts the use of hypothetical SQL statements to queries at the top-level, a forthcoming release will relax this restriction and allows to both create views with `assume` and also nested uses of assumptions.

4.6 Hypothetical Datalog Queries and Rules

A recent and novel addition to DES has been hypothetical Datalog queries and rules [SP13b], which is based on an intuitionistic semantics [NSS08, McC88, MG12]. As in the former section, extensional as well as intensional data can be assumed and, in addition, hypothetical rules are allowed. Hence, in a forthcoming release it is expected to translate hypothetical SQL queries into Datalog queries, and SQL hypothetical views to Datalog predicates. The syntax of a hypothetical literal is as follows:

```
Rule1 /\ ... /\ RuleN => Goal
```

which means that, assuming that the current database is augmented with the rules Rule_i ($1 \leq i \leq N$), then Goal is computed with respect to the current database which is augmented with these rules. Such query is also understood as a literal in the context of a rule, so that any rule can contain hypothetical goals, as in $a :- b \Rightarrow c$. In turn, any Rule_i can contain hypothetical goals. Variables in Rule_i are local to Rule_i (i.e., they are neither shared with other rules nor the goal). Moreover, a hypothetical literal does neither share variables with other literals nor the head of the rule in which it occurs. Also, assumed rules must be safe (in the sense of the safety conditions in [Ull88]). Borrowing an example from [Bon88], the question “Which are the students (s) which would be eligible to graduate ($\text{grad}/1$) if taking ($\text{take}/2$) db and lp was enough to graduate?” can be issued as:

```
(grad(S) :- take(S,db), take(S,lp)) => grad(S')
```

This query amounts to assume a new rule only in the context of solving the goal $\text{grad}(S')$. Note that the syntax of this rule follows the syntax of Datalog rules. A logical equivalent query would be: $\exists S' \text{grad}(S') \Leftarrow \forall S ((\text{take}(S,db) \wedge \text{take}(S,lp)) \Rightarrow \text{grad}(S))$.

Although not available yet, the hypothetical removal of tuples from relations is expected for a next version, both in Datalog and SQL.

5 A Note on Performance

Although DES has not been designed under any performance directive, some numbers can be taken to analyze its behavior w.r.t. other systems. Here, we consider the logic programming system XSB 3.2, as an example of an efficient open-source, in-memory, tabled deductive database,

Tuples	XSB	DES	DB2	P-DES	Result
200	197	135	360	816	40,000
400	283	557	1,459	3,035	160,000
600	416	1,266	3,270	6,740	360,000
800	572	2,287	5,783	11,912	640,000
1,000	768	3,646	9,100	18,590	1,000,000

Table 1: Performance Data

and the system IBM DB2 10.1.0, as a commercial, persistent, relational database system. These systems are compared with the DES in-memory system and also with the persistent approach as described in Section 3.3. In this last case, DB2 is the target system providing persistence via an ODBC connection.

As benchmark, we consider the cost of computing the Cartesian product of a table with itself, depending on the table size (benchmark instances). We focus therefore on retrieving to the main memory the result but without actually displaying the result in order to elide the display time. Each test has been run 10 times, the maximum and the minimum numbers have been discarded, and then the average has been computed. All benchmarks are given in milliseconds and have been run on an Intel Core2 Quad CPU at 2.4GHz and 3GB RAM, running Windows XP 32bit SP3. DB2 was accessed through the IBM DB2 ODBC driver version 10.01.00.872 and with its default configuration, as XSB.

Table 1 collects such time data expressed as milliseconds. Also, the benchmark instance is denoted in column Tuples, which shows the number of tuples in the table. The column Result shows the number of tuples in the result set which will be retrieved to main memory. Columns XSB, DES, DB2, and P-DES show the computation times for retrieving into main memory the result of the Cartesian product for XSB, DES for the in-memory database, DB2, and DES with the table as a persistent predicate mapped to DB2.

Figure 2 graphically collects these numbers for the same range of tuples as in the table, and shown in the horizontal axis. In turn, vertical axis shows computation times expressed in milliseconds. The number of computed tuples are also represented (which has been scaled down –50 times– to compare with the time to compute them) and it is depicted in the top curve. As illustrated, the ranking of computing times are in this order, from the best to the worst: XSB, DES, DB2, and P-DES.

A first observation from the figures is that the worst case corresponds to DES with persistent predicates, and its curve follows a similar slope to the number of computed tuples, meaning that the time complexity is linear in this number. A second observation is that the curve for the in-memory DES system performs better than DB2. This is clearly expected as DB2 takes more tasks than DES to compute such tuples. For instance, first, tables are not in-memory data structures, but otherwise rely on the OS file system. Nevertheless, it features an efficient buffer manager that tries to keep all data in main memory. And, second, it has to deal with transactional behavior, although with no logging as there are no updates in the tests. Finally, XSB achieves the best numbers as it is a C implementation which compiles the predicates and indexes tabled results with tries.

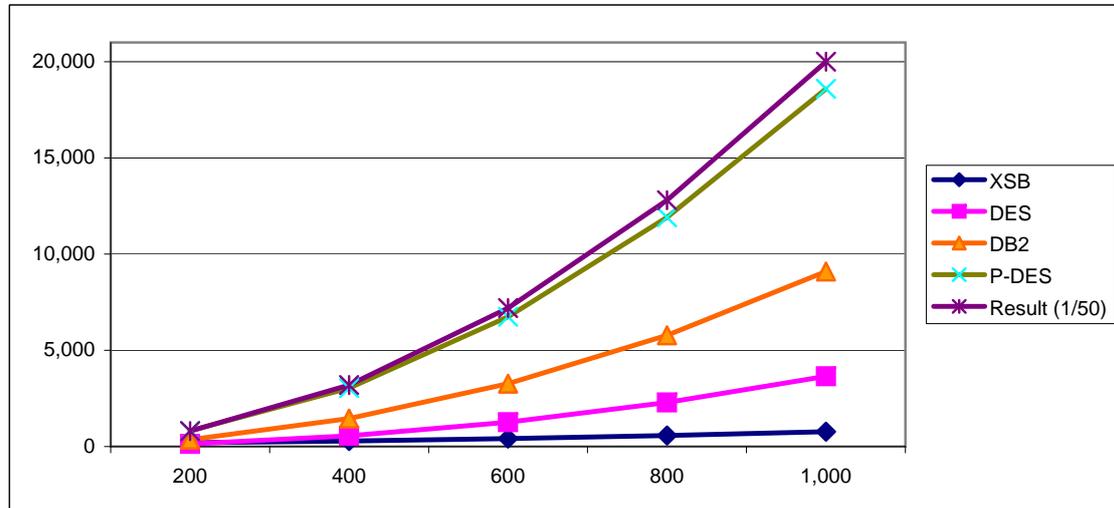


Figure 2: Performance Curves

Therefore, given these numbers, it can be expected to develop an efficient deductive system, in particular as fast as DB2, for solving these tests because the computation time due to the in-memory XSB system is negligible with respect to the persistence requirements for retrieving only 1,000 tuples from the external DBMS. This would be needed to read the table before actually computing the Cartesian product.

6 Conclusions

This paper has shown an ongoing educational project dealing with SQL expressiveness issues and allowing to project computations to external DBMS's. By providing a less-limited parser for practical expressiveness and projecting results from deductive databases for theoretical expressiveness, a better SQL language has been achieved. Although this is not the single system providing some of the expressive features presented, it is the only one which includes altogether all the ones presented in this paper. In addition, no other deductive system includes outer joins, intensional rule duplicates, SQL functional dependencies, the division operator, and hypothetical queries and rules, where some of these features are reported here for the first time (other interesting features include SQL algorithmic debugging, which are out of the scope of this paper). However, as an educational tool, it cannot compete with state-of-the-art systems w.r.t. performance. So, in addition to the enhancements mentioned along the paper, one of the obvious steps to continue this work is to apply known techniques to enhance performance (semi-naïve method for recursive queries [UI188], tabling trie-based indexing [SW12], native compilation, ...) and even to target to efficient tabled systems as XSB.

Bibliography

- [AOT⁺03] F. Arni, K. Ong, S. Tsur, H. Wang, C. Zaniolo. The Deductive Database System LDL++. *TPLP* 3(1):61–94, 2003.
- [Bon88] A. Bonner. Hypothetical Datalog: Complexity and Expressibility. *Theoretical Computer Science* 76:144–160, 1988.
- [CGC⁺04] J. Correas, J. M. Gómez, M. Carro, D. Cabeza, M. V. Hermenegildo. A Generic Persistence Model for (C)LP Systems (and Two Useful Implementations). In Jayaraman (ed.), *PADL*. LNCS 3057, pp. 104–119. Springer, 2004.
- [CGS08] R. Caballero, Y. García-Ruiz, F. Sáenz-Pérez. A Theoretical Framework for the Declarative Debugging of Datalog Programs. In *International Workshop on Semantics in Data and Knowledge Bases*. LNCS 4925, pp. 143–159. Springer, 2008.
- [CGS10] R. Caballero, Y. García-Ruiz, F. Sáenz-Pérez. Applying Constraint Logic Programming to SQL Test Case Generation. In *International Symposium on Functional and Logic Programming (FLOPS'10)*. LNCS 6009, pp. 191–206. Springer, 2010.
- [CGS11] R. Caballero, Y. García-Ruiz, F. Sáenz-Pérez. Algorithmic Debugging of SQL Views. In *Eighth Ershov Informatics Conference, PSI'11*. Pp. 204–210. June 2011.
- [CGS12] R. Caballero, Y. García-Ruiz, F. Sáenz-Pérez. Declarative Debugging of Wrong and Missing Answers for SQL Views. In *Eleventh International Symposium on Functional and Logic Programming (FLOPS'12)*. LNCS 7294, pp. 73–87. Springer, 2012.
- [Cod72] E. Codd. Relational Completeness of Data Base Sublanguages. In Rustin (ed.), *Database Systems*. Courant Computer Science Symposia Series 6. Englewood Cliffs, N.J. Prentice-Hall, 1972.
- [con06] Laws for Rewriting Queries Containing Division Operators. In *ICDE*. IEEE Computer Society, 2006.
- [Dat09] C. J. Date. *SQL and relational theory: how to write accurate SQL code*. O'Reilly, Sebastopol, CA, 2009.
- [Die] S. W. Dietrich. Extension Tables: Memo Relations in Logic Programming. In *IEEE Symp. on Logic Programming*. IEEE Computer Society.
- [Die01] S. Dietrich. *Understanding Relational Database Query Languages*. Prentice Hall, 2001.
- [DS92] C. Draxler, U. M. C. für Informations-und Sprachverarbeitung. *A Powerful Prolog to SQL Compiler*. CIS-Bericht-. Universität München, Centrum für Informations-und Sprachverarbeitung, 1992.
- [Got12] G. Gottlob. Second Datalog 2.0 Workshop. Austria, Vienna, 2012.

- [ISO00] ISO/IEC. ISO/IEC 132111-2: Prolog Standard. 2000.
- [ISO08] ISO/IEC. SQL:2008 ISO/IEC 9075(1-4,9-11,13,14):2008 Standard. 2008.
- [McC88] L. T. McCarty. Clausal Intuitionistic Logic I - Fixed-Point Semantics. *Journal of Logic Programming* 5(1):1–31, 1988.
- [McC03] L. I. McCann. On Making Relational Division Comprehensible. In *ASEE/IEEE Frontiers in Education Conference*. 2003.
- [MG12] D. Miller, N. Gopalan. *Programming with Higher-Order Logic*. Cambridge University Press, 2012.
- [NSS08] S. Nieva, F. Sáenz-Pérez, J. Sánchez. Formalizing a Constraint Deductive Database Language based on Hereditary Harrop Formulas with Negation. In *FLOPS'08, Proceedings*. LNCS 4989, pp. 289–304. Springer, 2008.
- [RU93] R. Ramakrishnan, J. Ullman. A survey of research on Deductive Databases. *Journal of Logic Programming* 23(2):125–149, 1993.
- [SP07] F. Sáenz-Pérez. ACIDE: An Integrated Development Environment Configurable for LaTeX. *The PracTeX Journal* 2007(3), August 2007. <http://acide.sourceforge.net>.
- [SP12] F. Sáenz-Pérez. Outer Joins in a Deductive Database System. *Electronic Notes in Theoretical Computer Science* 282(0):73 – 88, 2012.
- [SP13a] F. Sáenz-Pérez. Datalog Educational System 3.4. December 2013. <http://des.sourceforge.net/>.
- [SP13b] F. Sáenz-Pérez. Implementing Tabled Hypothetical Datalog. In *Proceedings of the 25th IEEE International Conference on Tools with Artificial Intelligence, ICTAI'13*. 2013.
- [SP12] F. Sáenz-Pérez. Tabling with Support for Relational Features in a Deductive Database. *ECEASST* 55, 2012.
- [SCG11] F. Sáenz-Pérez, R. Caballero, Y. García-Ruiz. A Deductive Database with Datalog and SQL Query Languages. In Yang (ed.), *APLAS*. LNCS 7078, pp. 66–73. Springer, 2011.
- [SK80] M. Stonebraker, K. Keller. Embedding Expert Knowledge and Hypothetical Data Bases into a Data Base System. In *Proceedings of the 1980 ACM SIGMOD International Conference on Management of Data*. SIGMOD '80, pp. 58–66. ACM, 1980.
- [SW12] T. Swift, D. S. Warren. XSB: Extending Prolog with Tabled Logic Programming. *TPLP* 12(1-2):157–187, 2012.
- [TLLP08] G. Terracina, N. Leone, V. Lio, C. Panetta. Experimenting with recursive queries in database and logic programming systems. *TPLP* 8(2):129–165, 2008.

- [Ull88] J. D. Ullman. *Principles of Database and Knowledge-Base Systems, Volume II*. Computer Science Press, 1988.
- [ZCF⁺97] C. Zaniolo, S. Ceri, C. Faloutsos, R. T. Snodgrass, V. S. Subrahmanian, R. Zicari. *Advanced Database Systems*. Morgan Kaufmann, 1997.