

# **An Introduction to Deductive Database Languages and Systems**

**Kotagiri Ramamohanarao and James Harland**

## **1. Introduction**

One of the most fundamental uses of a computer is to store and retrieve information, particularly when there are a large amount of data to be stored, or there are complex manipulations that must be performed on them. There has been a large amount of research on the most efficient techniques to store and retrieve data, and the associated problems now have satisfactory solutions. However, the problem of understanding and interpreting this large amount of information remains, particularly when the amounts of data belong to complex domains, such as those involving mineral exploration and financial analysis.

To tackle this problem, a mechanism for reasoning about the stored information is necessary. Such a mechanism must be able to cope with large amounts of information, as well as to perform sophisticated inferences, and to draw the appropriate conclusions. A framework in which these problems may be attacked is given by the field of deductive databases. Deductive databases not only store explicit information in the manner of a relational database, but they also store rules that enable inferences based on the stored data to be made. This area is an outgrowth of the field of logic programming, in which mathematical logic is used to directly model computational concepts. Together with techniques developed for relational databases, this basis in logic means that deductive databases are capable of handling large amounts of information as well as performing reasoning based on that information.

There are many application areas for deductive database technology. One area is that of decision support systems. In particular, the exploitation of an organization's resources requires not only sufficient information about the current and future status of the resources themselves, but also a way of reasoning effectively about plans for the future. The present generation of decision support systems are severely deficient when it comes to reasoning about future plans. Deductive database technology is an appropriate solution to this problem.

Another fruitful application area is that of expert systems. There are many computing applications in which there are large amounts of information, from which the important facts may be distilled by a simple yet tedious analysis. For example, medical analysis and monitoring can generate a large amount of data, and an error can have disastrous consequences. A tool to carefully monitor a patient's condition or to retrieve relevant cases during diagnosis reduces the risk of error in such

circumstances. Deductive database technology allows the analysis of these data to be performed more efficiently and with a lower chance of error than by ad hoc methods. Such an intelligent tool allows the human experts to concentrate on the main problems, rather than being distracted by details. A similar example may be found in mineral exploration; a large amount of data may be generated, which can then be analyzed for clues suggesting the presence of the desired mineral.

Planning systems are another application area. For example, a student planning a course of study at a university, or a passenger planning a round-the-world trip often need to consider a large body of information, as well as the ability to explore alternatives and hypotheses. A deductive database is able to advise students about pre-requisites and regulations on the choice of subjects, or a traveller of the financial implications of a given change in itinerary.

Deductive database systems have been the subject of extensive research, and several prototype deductive database systems are now emerging, as evidenced by the descriptions appearing elsewhere in this issue.

The rest of this introduction is organized as follows. In Section 2, we discuss various language issues for deductive database systems, and in Section 3 we describe implementation schemes for these systems. In Section 4, we briefly describe various implementations of deductive database systems, and in Section 5 we present our conclusions.

## 2. Deductive Database Languages

In this section we briefly discuss some language issues relevant to deductive databases. For more details, the reader is referred to Lloyd (1987).

The deductive database field has had close links with the logic programming community, and much of the development of deductive database systems has centered around languages based on Horn clauses. This class of formulas forms the basis of Prolog, and is powerful enough to encode Turing machines (Tärnlund, 1977).

A *Horn clause* is generally written as

$$p(\tilde{t}) :- q_1(\tilde{t}_1), \dots, q_n(\tilde{t}_n)$$

where  $p$  and  $q_1, \dots, q_n$  are predicate letters,  $n \geq 0$ , and all variables that occur in the terms  $\tilde{t}, \tilde{t}_1, \dots, \tilde{t}_n$  are considered universally quantified at the front of the clause.

Note that  $n$  may be 0, in which case we refer to the clause as a *fact*. Otherwise, we refer to the clause as a *rule*.

The atom  $p(\tilde{t})$  is referred to as the *head* of the clause, and  $q_1(\tilde{t}_1), \dots, q_n(\tilde{t}_n)$  as the *body* of the clause.

A *logic program* is a set of Horn clauses. The terms  $\tilde{t}, \tilde{t}_1, \dots, \tilde{t}_n$  may, in general, be arbitrary (first-order) terms, and hence may contain variables and/or function symbols.

It is often useful to consider sub-classes of this class of programs. A common restriction is to only allow terms to be either variables or constants. Such programs are known as *Datalog* programs. An important property of such programs is that it is decidable whether a given query is logically entailed by a Datalog program. Hence, it is reasonable to expect that a deductive database system should terminate on all Datalog programs.

Not all deductive database systems restrict programs to be Datalog programs. Datalog programs are somewhat restrictive; for example, the append program is not a Datalog program, as it requires the use of function symbols.

In the deductive database field, a distinction is usually made between predicates defined by rules alone (referred to as the *intentional* database or IDB), and predicates defined by facts alone (referred to as the *extensional* database or EDB). Any logic program can be rewritten so that all predicates are either IDB or EDB predicates. Often it is useful to consider a given IDB for various EDBs.

While Horn clauses are Turing complete (Tärnlund, 1977), it is common to extend the language of Horn clauses so that the body of a clause is a conjunction of *literals* (i.e., an atom or the negation of an atom, rather than a conjunction of atoms alone). The negative literals are inferred by the use of the *Negation as Failure* rule (Clark, 1978); a literal  $\neg A$  succeeds if  $A$  fails. The addition of this feature gives the language more expressive power, but it can also confuse the semantics of the program somewhat. For example, consider the program below.

```
p :- ¬ q
q :- ¬ p
```

Here it is not clear whether we should interpret  $p$  as true (and  $q$  as false, or vice-versa). As a result, negation generally has to be used carefully in logic programs to avoid problems of this kind. There has been a great deal of work on the semantics of negation in logic programs, and we give only a brief overview here. (For more information, see Gelfond and Lifschitz, 1988; Przymusiński, 1988; Van Gelder et al., 1991; Kemp et al., 1992).

A useful class of programs in which the use of negation is restricted is known as the class of *stratified* programs (Chandra and Harel, 1985; Blair and Walker, 1988). Intuitively, a program is stratified if there is no recursion through negation. For example, the following program, which defines the acyclic part of a graph, is stratified.

```
path(X,Y) :- edge(X,Y).
path(X,Y) :- edge(X,Z), path(Z,Y).
```

$$\text{acyclic}(X,Y) \text{ :- path}(X,Y), \neg \text{path}(Y,X).$$

Note that the definition of the *acyclic* predicate depends on the *path* predicate, but not vice-versa.

A more general class of programs, which is based on the same general idea, is the class of *locally stratified* programs (Przymusinski, 1988). Essentially, a locally stratified program allows recursion through negation, provided that no atom depends on its own negation. Further extensions of this concept include modular stratification (Ross, 1990).

Another restriction that is often imposed is to consider only programs that are *range-restricted*. A program is range-restricted if every variable that appears in the head of a clause also appears in the body of the clause (Bancilhon and Ramakrishnan, 1988) (note that this definition can be simply extended when negative literals appear in the body of a clause). This implies that all facts in the program must be ground (i.e., contain no variables). The main advantage of this class of programs is that in the query evaluation process, only matching is needed, which is significantly more efficient than full unification. Also, all answers to a given query are ground, and hence there is no need to check for answers subsuming one another.

Several of the prototype systems described in this article have implemented various combinations of the above features. Some systems only support Datalog range-restricted programs with stratified negation, some support modularly stratified programs, and/or function symbols. Some systems also do not impose any restrictions other than modular stratification. More details are provided in Section 4.

Many deductive database systems also include aggregate operators, such as sum, max, min, and count. While these operators allow the simple expression of many database programs, it is possible to write simple programs with a complicated semantics (as in the case for negations), and so many of the concepts introduced for negation (e.g., stratification) are also used for aggregate operators.

### 3. Implementation Schemes

There has been a significant body of research in the area of implementation of logic programming systems and deductive database systems, and a substantial body of theoretical work has been developed for such systems. In this article we are interested only in implementation techniques for deductive databases. These implementation techniques can be broadly categorized into three main groups:

- Prolog systems loosely coupled to database systems
- Top-down evaluation with memoing
- Bottom-up methods

### 3.1 Prolog Systems Loosely Coupled to Database Systems

Some of the early attempts to implement deductive databases were to interface a Prolog system to a database system (or a file store). These systems we refer to as Prolog database systems. These systems use Prolog computation and access appropriate database relations on a tuple-at-a-time basis. The benefit of this approach is that these systems can be implemented quickly and easily. The drawback of this approach is that the resulting system can be extremely inefficient because access to the underlying database system is tuple-at-a-time, and the resultant computation performed is similar to the nested loop join algorithm, but performing on several relations simultaneously. Several systems have been developed using this approach (Ramamohanarao et al., 1987; Zobel and Ramamohanarao, 1986; Horsfield et al., 1989).

Prolog is based on the top-down computation method, which is also known as backward chaining or SLD-resolution. This method is also used in theorem proving. It starts at the query and applies the rules of the program until it arrives at the facts.

The main steps in SLD-resolution are as follows:

1. Initialize the *goal list* of literals to the query.
2. Choose a goal  $A_i$  from the goal list  $A_1, A_2, \dots, A_i \dots A_n$ . Find a rule  $A :- B_1, \dots, B_m$  such that  $A\theta = A_i\theta$  for some most general unifier  $\theta$ . Terminate with failure if there are no such rules.
3. Update the goal list to  $(A_1, A_2, \dots, A_{i-1}, B_1, \dots, B_m, A_{i+1}, \dots, A_n)\theta$ .
4. If the goal list is not empty, go back to step 2. Otherwise, terminate with success; an answer to the query is contained in the substitutions.

Step 2 of the top-down algorithm has two forms of nondeterminism.

- The *computation rule* specifies which literal is to be selected.
- The *search rule* specifies the order in which the matching rules of the program are unified against the selected literal.

These two rules give the shape of the tree explored by the top-down algorithm.

In Prolog, the computation rule is to always use the leftmost literal in the goal list; the selected literal is replaced in the goal list by the body of the matching rule. The search rule is to always use the first matching rule.

Using this approach for deductive databases can result in a bottleneck, as large amounts of information can tend to “clog” the tuple-at-a-time nature of the computation. A significant development of this approach was pursued in the MegaLog system, developed at ECRC (Horsfield et al., 1989; Bocca, 1991).

MegaLog was designed to be similar to Prolog, but the main emphasis is on efficient database access. For example, MegaLog supports relational operations and indexing structures such as BANG files (Freeston, 1989).

Note that in such systems it is possible for some Datalog programs not to terminate, and hence it is the programmer's responsibility to ensure that all queries terminate.

### 3.2 Top-Down With Memoing

To overcome the problem of the termination of top-down methods on Datalog programs, the technique of *memoing* is often used. The main problem for termination of SLD-resolution is that the refutation procedure does not recognize goals that it has previously called, and so may loop needlessly. Methods of incorporating such a check into the SLD-resolution procedure have been studied by many researchers (Vieille, 1986, 1987, 1988; Dietrich, 1987; Warren, 1992), all of which may be considered variants of OLDT-resolution (Tamaki and Sato, 1986).

In its simplest form, top-down evaluation with memoing builds a tree similar to an SLD-tree except for the following restrictions and differences:

- Answers to subgoals are tabled (memoized) for future use: when the derivation proceeds from the goal list

$$A_1, \dots, A_i, \dots, A_n$$

to a descendent goal list of the form

$$(A_1, \dots, A_{i-1}, A_{i+1}, \dots, A_n)\theta$$

the atom  $A_i\theta$  is called an answer.

- If the subgoal  $A$  is an instance—possibly more instantiated—of a subgoal that occurred earlier in a left-to-right pre-order traversal of the tree, then  $A$  is not resolved using rules from the program, but is resolved against tabled answers.
- When a new answer is found, any subgoal that has been resolved using answers must be tested to see if it unifies with the new answer.

Although the above description is tuple-at-a-time in nature, it has been further developed to compute answers in an efficient set-at-a-time manner (Vieille, 1988).

Essentially, the search procedure “remembers” each goal that it has called, so that the evaluation of a given goal does not repeatedly derive the same subgoal. This system is in some respects a mixture of top-down and bottom-up methods, as many of the characteristics of the system have direct counterparts in bottom-up evaluation methods.

### 3.3 Bottom-Up Methods

The bottom-up method is known also as forward chaining or fixpoint computation. It starts at the facts and applies the rules until it arrives at the query. This approach is often used in the study of the semantics of logic programs, and by many deductive databases.

This computation method can be characterized by the following steps: Let the query be  $q(\tilde{Q})$ .

1. Initialize  $M$ , the set of known facts, to the set of facts in the program, and add the following rule to the program:

$$\text{ans}(\tilde{Q}) \text{ :- } q(\tilde{Q}).$$

2. For each rule  $A \text{ :- } A_1, \dots, A_n$ , look for substitutions  $\theta$  for which  $A_1\theta, \dots, A_n\theta \in M$ . For each such substitution, add  $A\theta$  to  $M$ .
3. If the set of known facts  $M$  has increased, go back to step 2.
4. The answer to the query is the set of ans facts in  $M$ .

The bottom-up approach naturally lends itself to the application of relational algebra techniques, as the conjunction of literals in the goal may be implemented by a sequence of join operations, for which many optimization techniques are known. However, the bottom-up method as described above completely ignores the values of any constants in a query, and therefore also derives facts which are irrelevant to the query. Relevant facts (including derived facts) are those which are used in the generation of answers to the query. The number of these irrelevant facts can be very large, and in general this can make bottom-up computation very expensive.

By contrast, the top-down method, with or without memoing, does not have this problem, because query evaluation uses the instantiated variables of the goal. To make bottom-up methods concentrate only on facts that are relevant to the query, techniques such as *magic sets* have been developed (Bancilhon et al., 1986; Beeri and Ramakrishnan, 1987, Balbin et al., 1991). This is one of the most important optimization techniques for bottom-up methods.

This is a source-to-source transformation; it transforms the *program* (the rules of the database) into another program that can be evaluated more efficiently by the standard bottom-up computation we have presented.

The magic set transformation is a general transformation, and can be applied to all programs. However, special care is needed when dealing with programs containing negations. The transformation provides a focus equivalent to top-down computation, so that only facts relevant to the query are generated.

For example, consider the program and query below.

?- partof(2, Y).

```
partof(X,Y) :- component(X,Y).
partof(X,Y) :- component(X,Z), partof(Z,Y).
```

Under the magic set transformation, the program and query become

```
?- partof_m(2, Y).
magic_partof(2).
magic_partof(Z) :- magic_partof(X), component(X,Z).
partof_m(X,Y) :- magic_partof(X), component(X,Y).
partof_m(X,Y) :- magic_partof(X), component(X,Z),
    partof_m(Z,Y).
```

The standard bottom-up evaluation of these rules produces the same result for this query as the evaluation of the unmodified rules would, but it looks at only the relevant facts. The `magic_partof` relation initially contains only the tuple  $\langle 2 \rangle$ , the input value for the first argument of `partof`. At each stage in the bottom-up evaluation of `magic_partof`, the computation adds to this relation the values of the first argument of `partof` that a top-down evaluation of the query would see at the corresponding depth in the search tree. At the end, `magic_partof` contains the *magic set* (i.e., all the values for the first argument of *partof* that the top-down evaluation of the query would ever see).

The modified rules of *partof* then use the magic set to avoid computing the parts of the *partof* relation that are not relevant to the query.

Below we provide a very simple form of the magic set transformation algorithm, which may be applied to any program.

- For each derived predicate of the program, create a magic predicate by prefixing `magic_` to the predicate name. The arguments of this new predicate are the bound arguments of the original predicate.
- For each rule, add a magic atom to the front of the rule; the arguments of this atom are the bound arguments of the rule head.
- For each modified rule of the program, create a new rule for each call to a derived predicate `p` whose bound arguments are  $X_1, X_2, \dots, X_n$ . The head of this new rule is `magic_p(X1, X2, ..., Xn)` and the body is the literals preceding the call.

There are several other optimization algorithms which have been developed, some for particular classes of programs such as linear recursions (Kemp et al., 1990; Harland and Ramamohanarao, 1993), and various others which are more generally applicable (Sacca and Zaniolo, 1986, 1987; Beerli and Ramakrishnan, 1987; Sagiv, 1990; Harland and Ramamohanarao, 1992; Kemp and Stuckey, 1993).

Several of the prototype systems described below have implemented various combinations of the above methods and techniques. Most systems implement either bottom-up evaluation or top-down with memoing, and various combinations of optimization techniques. Some allow the user control over the optimizations, while others select optimization strategies automatically. More details are provided in the next section.

#### 4. Prototype Deductive Database Systems

A large number of prototype deductive database systems have been developed to date. Several of the implemented systems are memory-based. These systems assume that all the required permanent relations can be kept in main memory, and during the process of computation, any temporary relations generated can also be kept in memory. Although this method suffices for applications where the temporary relations generated are small enough to fit into main memory, this is an unreasonable expectation for some applications. When this assumption is false, these systems tend to behave poorly; therefore techniques used in building relational database systems must be used.

Several of the implementations also assume that there is a single user of the database, and in general do not support transaction processing and crash recovery. In addition, many systems do not support essential database features such as integrity constraints and triggers. In spite of these limitations, substantial progress has been made towards demonstrating the feasibility of deductive database technology, and some prototype systems have been developed that do provide the expected features of a traditional database system. There are also commercial database systems under development that have the capabilities of a deductive database.

Below we give an overview of the state of development of various prototype systems. This overview is not a complete survey of all efforts that have taken place in the development of deductive database systems. We concentrate on systems which have had significant developmental effort and have received significant attention in the literature. We refer interested readers to a forthcoming survey article (Ramakrishnan and Ullman, in press), which covers other issues.

**RDL/C.** RDL/C is a programming language developed to integrate a rule-based language and the programming language C. RDL/C is derived from RDL1 (Maindreville and Simon, 1988). This language supports rules and abstract data types (Gardarin et al., 1989); therefore, the user can program at a higher level than is possible using the combination of SQL and C. In particular, the user does not have to manage temporary relations, which are handled by the system. Programs written in RDL/C are compiled into an embedded database query language. This approach has the advantage of being easy to integrate

into an object-oriented database system or a relational database system; this provides a powerful and flexible database system. In many respects, this system is similar to LOLA (see below).

**MegaLog.** MegaLog was developed at the European Computer Research Center (ECRC) (Horsfield et al., 1989; Bocca, 1991). This system is designed to support the manipulation of large amounts of data while also providing standard Prolog features. One of the main contributions of this development is the support of a multi-dimensional grid file system called Balanced And Nested Grid file (BANG). Other important features include its support for garbage collection, and excellent facilities for dictionary management. Because the behavior of the system is similar to Prolog, it does not guarantee termination even for Datalog programs. However, the system has proved to be a good development platform for data-intensive knowledge bases, such as the EKS system described below.

**EKS.** The ECRC Knowledge base System (EKS) was developed at ECRC from 1989 to 1991 (Vieille et al., 1990). Like several deductive database systems, one of the goals of this project is to demonstrate the viability of deductive database technology for real-world applications. EKS is built on the MegaLog Prolog platform (Horsfield et al., 1989; Bocca, 1991). The language of EKS is Datalog (and hence does not support function symbols). The main features of the EKS system include support for a very general form of integrity constraints, which may include references to recursive predicates and aggregate operations, rules which may contain recursion through aggregates, support for materialized views, and support for hypothetical query facilities. In this system, support for procedural definitions and updates is provided by the underlying MegaLog platform. The initial system was a single-user system. The computational model used in this system is derived from Query/SubQuery evaluation, a set-oriented top-down evaluation scheme with memoing (Vieille, 1986, 1989; Lefebvre and Vieille, 1989). One of the main advantages of this approach is that negation is handled in a top-down setting. This makes negation simpler to implement than bottom-up methods using the magic set transformation.

**LDL.** The LDL system was developed at Micro Computer Corporation (Naqvi and Tsur, 1989). One of the main features of this system is support for sets in the language. This system was built based on the bottom-up computation model, and uses several optimization techniques, such as magic sets. This system is a single user system, and all relations are memory-resident. The deductive part of this system is memory-resident. Later versions of the LDL system allow it to be interfaced to traditional relational systems, thus providing traditional database features such as transactions. A second-generation version of LDL,

known as LDL++ has been re-implemented (Zaniolo et al., 1993). Its main enhancements are the provision of interfaces to procedures written in C or C++, as well as the addition of abstract data types to the language.

**LOLA.** The LOLA system was developed at the Technical University of Munich (Freitag et al., 1991). The system is implemented by compiling Horn clauses, which may contain lists, into a Relational Lisp program with embedded SQL statements. The system does optimizations to minimize the calls to the underlying SQL database system. The system's support for multiple users and transactions is mainly derived from the underlying system. This implementation approach is very similar to that of Declare and SDS, described below (and in an article in this issue). The deductive part of this system is memory-resident, and hence this system is not scaleable to large databases when the intermediate relations are large.

**CORAL.** The CORAL system was developed at the University of Wisconsin at Madison (see article in this issue). CORAL uses bottom-up evaluation, with a wide variety of optimization strategies, which are specified by the programmer. One of the main features of CORAL is support for non-ground terms. The system is a single-user system and memory-resident. However the system can be connected to the EXODUS storage manager for access to permanent relations. It is not clear whether this kind of integration will scale up to large databases in performance terms.

**Glue-Nail.** The Glue-Nail system was developed at Stanford University (see article in this issue). An important feature of this system is the provision of two languages: one (Nail) for purely declarative statements based on Horn clauses, and another (Glue), which is procedural and used for I/O, updates and control constructs. The system also supports a form of higher-order syntax for the management of relation names. The system is a single-user system and memory-resident.

**Aditi.** The Aditi system was developed at the University of Melbourne (see article in this Issue). Aditi uses a bottom-up approach using relational technology. In this system, both permanent and temporary relations can be disk-resident, and hence the system is scaleable to large databases. The system supports function symbols, negation, and aggregates (including recursively defined aggregates). The architecture of this system is based on the client-server model, and supports parallel query processing. The system is also a multi-user system. Another important feature of Aditi is that bottom-up and top-down computations can be interleaved. The user can declare that a particular predicate is to be evaluated in a top-down fashion. Aditi then

makes a call to a Prolog system to execute such predicates. This mixing of top-down and bottom-up computation can improve performance by several orders of magnitude. However, for such predicates it is the responsibility of the user to ensure termination.

**Declare and SDS.** The Declare and SDS project is one of the earliest deductive database projects to build a commercial deductive database system (see article in this issue). This system has a lot of similarities to the LOLA system, although the Declare and SDS system is further developed. The language of this system is based on Horn clauses and supports lists, but with rules defining the same head predicate grouped together to form a virtual relation. The system is implemented using Relational Lisp, and is built on top of an extended version of the TransBase system. The system also provides support for types, as well as for distributed databases, and facilities for transactions.

**XSB.** The XSB system (Warren, 1989, 1992) was developed at Stony-Brook University. In many respects, this system has similar goals to the CORAL system in supporting non-ground terms and negation. However, the main distinction is the model of computation used in XSB, which is based on OLDT resolution, a top-down method with memoing (Tamaki and Sato, 1986; Dietrich, 1987). In this respect the XSB system resembles the EKS system. Like CORAL, this system is a single-user, memory-based system.

**Starburst.** The Starburst system (Haas et al., 1990) was developed at IBM Almaden. This is a substantial project, with the main goal being extensibility of the database system, and with some interest in deductive capabilities. The system supports a restricted but useful class of recursive rules. Due to this restriction, the system is able to use efficient specialized algorithms for query evaluation. The usefulness of the magic set transformation for non-recursive programs is demonstrated in this system (Mumick and Pirahesh, 1994).

**Commercial Systems.** In addition to Declare and SDS discussed above, there is also a commercial system currently under development at Groupe Bull. We believe that its main features include support for object-oriented features combined with the deductive facilities of EKS.

Some other interesting systems, such as ConceptBase (Jeusfeld and Staudt, 1993), COL (Abiteboul and Grumbach, 1991), LogicBase, Hy+ (Consens and Mendelzon, 1993), and X4 (Moerkotte and Lockemann, 1991), are discussed in the survey paper (Ramakrishnan and Ullman, in press).

## 5. Conclusion

Deductive database technology has now reached a level of maturity so that the commercial development of deductive database systems is feasible. There are several substantial prototype deductive database systems currently available from universities and research institutions, and it is now possible to build real applications using this technology. These prototype systems have already demonstrated the potential of deductive database systems to perform as efficiently as relational systems (for those applications where relational systems are appropriate). In addition, deductive database systems provide significantly more expressive power, both for querying the database and modeling of data.

However, before deductive database technology is generally accepted in the database community, these systems will need to have the standard database facilities for transaction processing, crash recovery, multi-user access, integrity constraints, triggers, and distributed database access. Unfortunately, several of the prototype systems do not have these facilities. We believe that systems such as Aditi, and Declare and SDS are closer to this goal than many others. It is also encouraging to see that there are some commercial deductive database systems under development which will include these standard database features.

## References

- Abiteboul, S. and Grumbach, S. A rule-based language with functions and sets. *ACM Transactions on Database Systems*, 16(1):1-30, 1991.
- Apt, K.R., Blair, H., and Walker, A. Towards a theory of declarative knowledge. In: Minker, J., ed., *Foundations of Deductive Databases and Logic Programming* Los Altos: Morgan Kaufmann, 1988, pp. 89-144.
- Balbin, I., Kemp. D., Meenakshi, K., and Ramamohanarao, K. Propagating constraints in recursive deductive databases. *Proceedings of the North American Conference on Logic Programming*, Cleveland, OH, 1989.
- Balbin, I., Port, G., Ramamohanarao, K., and Meenakshi, K. Efficient bottom-up computation of queries on stratified databases. *Journal of Logic Programming*, 11:295-345, 1991.
- Bancilhon, F., Maier, D., Sagiv, Y., and Ullman, J. Magic sets and other strange ways to implement logic programs. *Proceedings of the ACM Symposium on the Principles of Database Systems*, Cambridge, 1986.
- Bancilhon, F. and Ramakrishnan, R. Performance evaluation of data intensive logic programs/ In: Minker, J., ed., *Foundations of Deductive Databases and Logic Programming*, Los Altos: Morgan Kaufmann, 1988.
- Beeri, C. and Ramakrishnan, R. On the power of magic. *Proceedings of the ACM Symposium on the Principles of Database Systems*, San Diego, CA, 1987.

- Bocca, J. MegaLog: A platform for developing knowledge base management systems. *Proceedings of the Second International Symposium on Database Systems for Advanced Applications*, Tokyo, 1991.
- Chandra, A. and Harel, D. Horn clause queries and generalizations. *Journal of Logic Programming*, 2(1):1-15, 1985.
- Consens, M. and Mendelzon, A. Hy: A hygraph-based query and visualization system. *Proceedings of the ACM SIGMOD Annual Conference on Management of Data*, Washington, DC, 1993.
- Clark, K. Negation as failure. In: Gallaire, H. and Minker, J., eds. *Logic and Databases*, New York: Plenum Press, 1978.
- Dietrich, S. Extension tables: Memo relations in logic programming. *Proceedings of the Symposium on Logic Programming*, San Francisco, 1987.
- Freeston, M. Advances in the design of the BANG file. *Proceedings of the Third International Conference on the Foundations of Data Organization and Algorithms*, Paris, 1989.
- Freitag, B., Schütz, H., and Specht, G. LOLA: A logic language for deductive databases and its implementation. *Proceedings of the Second International Symposium for Advanced Applications*, Tokyo, 1991.
- Gardarin, G., Cheiney, J., Kiernan, G., Pastre, D., and Stora, H. Managing complex objects in an extensible relational DBMS. *Proceedings of the Fifteenth International Conference on Very Large Databases*, Amsterdam, 1989.
- Gelfond, M. and Lifschitz, V. The stable model semantics for logic programming. *Proceedings of the Fifth International Conference and Symposium on Logic Programming*, Seattle, 1988.
- Haas, L., Chang, W., Lohman, G., McPherson, J., Wilms, P., Lapis, G., Lindsay, B., Pirahesh, H., Carey, M., and Shekita, E. Starburst mid-flight: As the dust clears. *IEEE Transactions on Knowledge and Data Engineering*, 2(1):143-160, 1990.
- Harland, J. and Ramamohanarao, K. Constraints for query optimisations in deductive databases. *Proceedings of the Second Far East Workshop on Future Database Systems*, Kyoto, 1992.
- Harland, J. and Ramamohanarao, K. Constraint propagation for linear recursive rules. *Proceedings of the International Conference on Logic Programming*, Budapest, 1993.
- Horsfield, T., Bocca, J., and Dahmen, M. MegaLog User Guide, Technical Report, ECRC, Munich, 1989.
- Jeusfeld, M. and Staudt, M. Query optimisation in deductive object bases. In: Vossen, G., Freytag, J., and Maier, D., eds., *Query Processing for Advanced Database Applications*, Los Altos: Morgan-Kaufmann, 1993.
- Kemp, D., Ramamohanarao, R., and Somogyi, Z. Right-, left-, and multi-linear rule transformations that maintain context information, *Proceedings of the Sixteenth*

- International Conference on Very Large Data Bases*, Brisbane, 1990.
- Kemp, D., Srivastava, D., and Stuckey, P. Query restricted bottom-up evaluation of normal programs. *Proceedings of the Joint International Conference and Symposium on Logic Programming*, Washington DC, 1992.
- Kemp, D. and Stuckey, P. Analysis based constraint query optimization. *Proceedings of the International Conference on Logic Programming*, Budapest, 1993.
- LeFebvre, A. and Vieille, L. On deductive query evaluation in the dedgin\* system. *Proceedings of the First International Conference on Deductive and Object-Oriented Databases*, Kyoto, 1989.
- Lloyd, J. *Foundations of Logic Programming*, (2nd ed.), Berlin: Springer-Verlag, 1987.
- de Maindreville, C. and Simon, E. A production rule based approach to deductive databases. *Proceedings of the Fourth IEEE International Conference on Data Engineering*, Los Angeles, 1988.
- Moerkotte, G. and Lockemann, P. Reactive consistency control in deductive databases, *ACM Transactions on Database Systems*, 16(4):670-702, 1991.
- Mumick, I. and Pirahesh, H. Implementation of magic-sets in Starburst. *Proceedings of the ACM SIGMOD International Conference on the Management of Data*, Minneapolis, 1994.
- Naqvi, S. and Tsur, S. *A Logical Language for Data and Knowledge Bases*. Computer Science Press, 1989.
- Przymusinski, T. On the declarative semantics of stratified deductive databases. In: Minker, J., ed., *Foundations of Deductive Databases and Logic Programming*, Los Altos: Morgan Kaufmann, 1988, 193–216.
- Ramakrishnan, R. and Ullman, J. A survey of research on deductive database systems. *Journal of Logic Programming*, in press.
- Ramamohanarao, K., Shepherd, J., Balbin, I., Port, G., Naish, L., Thom, J., Zobel, J., and Dart, P. The NU-Prolog deductive database system. *IEEE Data Engineering*, 10(4):10-19, 1987.
- Ross, K. Modular stratification and magic sets for datalog programs with negation. *Proceedings of the ACM Symposium on Principles of Database Systems*, Nashville, 1990.
- Sacca, D. and Zaniolo, C. The generalized counting method for recursive queries. *Proceedings of the First International Conference on Database Theory*, Rome, 1986.
- Sacca, D. and Zaniolo, C. Magic counting methods. *Proceedings of the ACM SIGMOD International Conference on Management of Data*, Nashville, 1987.
- Sagiv, Y. Is there anything better than magic? *Proceedings of the North American Conference on Logic Programming*, Austin, TX 1990.
- Tamaki, H. and Sato, T. OLD resolution with tabulation. *Proceedings of the Third International Conference on Logic Programming*, London, 1986.
- Tärnlund, S.-A. Horn clause computability. *BIT*, 17:215-226, 1977.

- Van Gelder, A., Ross, K., and Schlipf, J. The well-founded semantics for general logic programs. *Journal of the ACM*, 38(3):620-50, 1991.
- Vieille, L. Recursive axioms in deductive databases: The query-subquery approach. *Proceedings of the First International Conference on Expert Database Systems*, Charleston, SC, 1986.
- Vieille, L. Database complete proof procedures based on SLD-resolution. *Proceedings of the Fourth International Conference on Logic Programming*, Melbourne, Australia, 1987.
- Vieille, L. From QSQ towards QoSQ: Global optimization of recursive queries. *Proceedings of the Second International Conference on Expert Database Systems*, Tysons Corner, VA, 1988.
- Vieille, L. Recursive query processing: The power of logic. *Theoretical Computer Science*, 69(1):1-53, 1989.
- Vieille, L., Bayer, P., Küchenhoff, V., and LeFebvre, A. EKS-V1, a short overview. *AAAI-90 Workshop on Knowledge Based Management Systems*, 1990.
- Warren, D. The XWAML: A machine that integrates Prolog and deductive database query evaluation. Technical Report 89/25, Department of Computer Science, SUNY at Stony Brook, NY, October, 1989.
- Warren, D. Memoing for logic programs. *Communications of the ACM*, 35(3):93-111, 1992.
- Zaniolo, C., Arni, N., and Ong, K. Negation and aggregates in recursive rules: The LDL++ approach. *Proceedings of the International Conference on Deductive and Object-Oriented Databases*, Phoenix, AZ, 1993.
- Zobel, J. and Ramamohanarao, K. Accessing existing databases from Prolog technical report 86/17, Department of Computer Science, University of Melbourne, 1986.