

## *The deductive database system $\mathcal{L}\mathcal{D}\mathcal{L}++$*

FAIZ ARNI

*InferData Corporation, 8200 N. MoPac Expressway, Austin, TX 78759, USA*  
(e-mail: [farni@inferdata.com](mailto:farni@inferdata.com))

KAYLIANG ONG

*Trilogy Inc., 5001 Plaza on the Lake, Austin, TX 78746, USA*  
(e-mail: [kayliang.ong@trilogy.com](mailto:kayliang.ong@trilogy.com))

SHALOM TSUR

*BEA Systems, 2315 N. First Street, San Jose, CA 95131, USA*  
(e-mail: [dicktsur@bea.com](mailto:dicktsur@bea.com))

HAI XUN WANG

*IBM T. J. Watson Research Center, 30 Saw Mill River Rd., Hawthorne, NY 10532, USA*  
(e-mail: [haixun@us.ibm.com](mailto:haixun@us.ibm.com))

CARLO ZANIOLO

*Computer Science Department, University of California, Los Angeles, CA 90095, USA*  
(e-mail: [zaniolo@cs.ucla.edu](mailto:zaniolo@cs.ucla.edu))

---

### Abstract

This paper describes the  $\mathcal{L}\mathcal{D}\mathcal{L}++$  system and the research advances that have enabled its design and development. We begin by discussing the new nonmonotonic and nondeterministic constructs that extend the functionality of the  $\mathcal{L}\mathcal{D}\mathcal{L}++$  language, while preserving its model-theoretic and fixpoint semantics. Then, we describe the execution model and the open architecture designed to support these new constructs and to facilitate the integration with existing DBMSs and applications. Finally, we describe the lessons learned by using  $\mathcal{L}\mathcal{D}\mathcal{L}++$  on various tested applications, such as middleware and datamining.

**KEYWORDS:** deductive databases, logic and databases, Datalog, nonmonotonic reasoning, monotonic aggregates

---

### 1 Introduction

The  $\mathcal{L}\mathcal{D}\mathcal{L}++$  system, which was completed at UCLA in the summer of 2000, concludes a research project that was started at MCC in 1989 in response of the lessons learned from of its predecessor, the  $\mathcal{L}\mathcal{D}\mathcal{L}$  system. The  $\mathcal{L}\mathcal{D}\mathcal{L}$  system, which was completed 1988, featured many technical advances in language design (Naqvi and Tsur, 1989), and implementation techniques (Chimenti *et al.*, 1990). However, its deployment in actual applications (Tsur, 1990a, 1990b) revealed many problems and needed improvements, which motivated the design of the new  $\mathcal{L}\mathcal{D}\mathcal{L}++$  system. Many of these problems were addressed in the early versions of the  $\mathcal{L}\mathcal{D}\mathcal{L}++$  prototype that were built at MCC in the period 1990–1993; but other problems,

particularly limitations due to the stratification requirement, called for advances on nonmonotonic semantics, for which solutions were discovered and incorporated into the system over time – until the last version (Version 5.1) completed at UCLA in the summer of 2000.

In this paper, we concentrate on the most innovative and distinctive features of *LDL++*, which can be summarized as follows:

- Its new language constructs designed to extend the expressive power of the language, by allowing negation and aggregates in recursion, while retaining the declarative semantics of Horn clauses,
- Its execution model designed to support (i) the new language constructs, (ii) data-intensive applications via tight coupling with external databases, and (iii) an open architecture for extensibility to new application domains,
- Its extensive application testbed designed to evaluate the effectiveness of deductive database technology on data intensive applications and new domains, such as middleware and data mining.

## 2 The language

A challenging research objective pursued by *LDL++* was that of extending the expressive power of logic-based languages beyond that of *LDL* while retaining a fully declarative model-theoretic and fixpoint semantics. As many other deductive database systems designed in the 80s (Minker, 1996), the old *LDL* system required programs to be stratified with respect to nonmonotonic constructs such as negation and set aggregates (Ramakrishnan and Ullman, 1995). While stratification represented a major step forward in taming the difficult theoretical and practical problems posed by nonmonotonicity in logic programs, it soon became clear that it was too restrictive for many applications of practical importance. Stratification makes it impossible to support efficiently even basic applications, such as Bill of Materials and optimized graph-traversals, whose procedural algorithms express simple and useful generalizations of transitive closure computations. Thus, deductive database researchers have striven to go beyond stratification and allow negation and aggregates in the recursive definitions of new predicates. *LDL++* provides a comprehensive solution to this complex problem by the fully integrated notions of (i) choice, (ii) User Defined Aggregates (UDAs), and (iii) XY-stratification. Now, XY-stratification generalizes stratification to support negation and (nonmonotonic) aggregates in recursion. However, the choice construct (used to express functional dependency constraints) defines mappings that, albeit nondeterministic, are monotonic and can thus be used freely in recursion. Moreover, this construct makes it possible to provide a formal semantics to the notion of User-Defined Aggregates (UDAs), and to identify a special class of UDAs that are monotonic (Zaniolo and Wang, 1999); therefore, the *LDL++* compiler recognizes monotonic UDAs and allows their unrestricted usage in recursion. In summary, *LDL++* provides a two-prong solution to the nonmonotonicity problem, by (i) enlarging the class of logic-based constructs that are monotonic (with constructs such as choice and monotonic aggregates), and (ii) supporting XY-stratification for hard-core nonmonotonic constructs, such as negation and nonmonotonic aggregates.

These new constructs of  $\mathcal{L}\mathcal{D}\mathcal{L}++$  are fully integrated with all other constructs, and easy to learn and use. Indeed, a user needs not to know abstract semantic concepts, such as stable models or well-founded models; instead, the user only needs to follow simple syntactic rules – the same rules that are then checked by the compiler. In fact, the semantic well-formedness of  $\mathcal{L}\mathcal{D}\mathcal{L}++$  programs can be checked at compile time – a critical property of stratified programs that was lost in later extensions, such as modular stratification (Ross, 1994). These new constructs are described next.

### 2.1 Functional constraints

Say that we have a database containing the relations `student(Name, Major, Year)` and `professor(Name, Major)`. In fact, let us take a toy example that only has the following facts<sup>1</sup>

```
student('JimBlack', ee, senior).      professor(ohm, ee).
                                     professor(bell, ee).
```

Now, the rule is that the major of a student must match his/her advisor's major area of specialization. Then, eligible advisors can be computed as follows:

```
elig_adv(S,P) ← student(S, Majr, Year), professor(P, Majr).
```

This yields

```
elig_adv('JimBlack', ohm).
elig_adv('JimBlack', bell).
```

However, since a student can only have one advisor, the goal `choice((S), (P))` must be added to our rule to force the selection of a unique advisor for each student:

#### Example 2.1

*Computation of unique advisors by a choice rule*

```
actual_adv(S,P) ← student(S, Majr, Yr), professor(P, Majr),
                  choice((S), (P)).
```

The goal `choice((S), (P))` can also be viewed as enforcing a *functional dependency* (FD)  $S \rightarrow P$  on the results produced by the rule; thus, in `actual_adv`, the second column (professor name) is functionally dependent on the first one (student name). Therefore, we will refer to `S` and `P`, respectively, as the left side and the right side of this FD, and of the choice goal defining it. The right side of a choice goal cannot be empty, but its left side can be empty, denoting that all tuples produced must share the same values for the right side attributes.

The result of the rule of Example 2.1 is *nondeterministic*: it can either return a singleton relation containing the tuple ('JimBlack', ohm), or one containing the tuple ('JimBlack', bell).

<sup>1</sup> We follow the standard convention of using upper case initials to denote variables; lower case initials and strings enclosed in quotes denote constants.

A program where the rules contain choice goals is called a *choice program*. The semantics of a choice program  $P$  can be defined by transforming  $P$  into a program with negation,  $foe(P)$ , called the *first order equivalent* of  $P$ . Now,  $foe(P)$  exhibits a multiplicity of stable models, each obeying the FDs defined by the choice goals; each such stable model corresponds to an alternative set of answers for  $P$  and is called a *choice model* for  $P$ . The first order equivalent of Example 2.1 is as follows:

*Example 2.2*

The first order equivalent for Example 2.1

$$\begin{aligned} \text{actual\_adv}(S,P) &\leftarrow \text{student}(S,\text{Majr},\text{Yr}), \text{professor}(P,\text{Majr}), \\ &\quad \text{chosen}(S,P). \\ \text{chosen}(S,P) &\leftarrow \text{student}(S,\text{Majr},\text{Yr}), \text{professor}(P,\text{Majr}), \\ &\quad \neg \text{diffChoice}(S,P). \\ \text{diffChoice}(S,P) &\leftarrow \text{chosen}(S,P'), P \neq P'. \end{aligned}$$

This can be read as a statement that a professor will be assigned to a student whenever a different professor has not been assigned to the same student. In general,  $foe(P)$  is defined as follows:

*Definition 2.1*

Let  $P$  denote a program with choice rules: its first order equivalent  $foe(P)$  is obtained by the following transformation. Consider a choice rule  $r$  in  $P$ :

$$r : A \leftarrow B(Z), \text{choice}((X_1),(Y_1)), \dots, \text{choice}((X_k),(Y_k)).$$

while

- (i)  $B(Z)$  denotes the conjunction of all the goals of  $r$  that are not choice goals, and
- (ii)  $X_i, Y_i, Z, 1 \leq i \leq k$ , denote vectors of variables occurring in the body of  $r$  such that  $X_i \cap Y_i = \emptyset$  and  $X_i, Y_i \subseteq Z$ .

Then,  $foe(P)$  is constructed from  $P$  as follows:

1. Replace  $r$  with a rule  $r'$  obtained by substituting the choice goals with the atom  $\text{chosen}_r(W)$ :

$$r' : A \leftarrow B(Z), \text{chosen}_r(W)$$

where  $W \subseteq Z$  is the list of all variables appearing in choice goals, i.e.  $W = \bigcup_{1 \leq j \leq k} X_j \cup Y_j$ .

2. Add the new rule

$$\text{chosen}_r(W) \leftarrow B(Z), \neg \text{diffChoice}_r(W).$$

3. For each choice atom  $\text{choice}((X_i),(Y_i))$  ( $1 \leq i \leq k$ ), add the new rule

$$\text{diffChoice}_r(W) \leftarrow \text{chosen}_r(W'), Y_i \neq Y'_i$$

where (i) the list of variables  $W'$  is derived from  $W$  by replacing each  $A \notin X_i$  with a new variable  $A'$  (i.e. by priming those variables), and (ii)  $Y_i \neq Y'_i$  denotes

the inequality of the vectors, i.e.  $Y_i \neq Y'_i$  is true when for some variable  $A \in Y_i$  and its primed counterpart  $A' \in Y'_i$ ,  $A \neq A'$ .

#### Monotonic nondeterminism

##### Theorem 2.1

Let  $P$  be a positive program with choice rules. Then the following properties hold (Giannotti *et al.*, 1991):

- $foe(P)$  has one or more total stable models.
- The *chosen* atoms in each stable model of  $foe(P)$  obey the FDs defined by the choice goals.

Observe that the  $foe(P)$  of a program with choice does not have total well-founded models; in fact, for our Example 2.1, the well-founded model yields undefined values for advisors. Therefore, the choice construct can express nondeterministic semantics, which can be also expressed by stable models, but not by well-founded models. On the other hand, the choice model avoids the exponential complexity which is normally encountered under stable model semantics.

Indeed, the computation of stable models is  $\mathcal{NP}$ -hard (Schlipf, 1993), but the computation of choice models for positive programs can be performed in polynomial time with respect to the size of the database. This, basically, is due to the monotonic nature of the choice construct that yields a simple fixpoint computation for programs with choice (Giannotti *et al.*, 2001b). Indeed, the use of choice rules in positive programs preserves their monotonic properties. A program  $P$  can be viewed as consisting of two separate components: an extensional component (i.e. the database facts), denoted  $edb(P)$ , and an intensional one (i.e. the rules), denoted  $idb(P)$ . Then, a positive choice program defines a monotonic multi-valued mapping from  $edb(P)$  to  $idb(P)$ , as per the following theorem proven in Giannotti *et al.* (2001b):

##### Theorem 2.2

Let  $P$  and  $P'$  be two positive choice programs where  $idb(P') = idb(P)$  and  $edb(P') \supseteq edb(P)$ . Then, if  $M$  is a choice model for  $P$ , then, there exists a choice model  $M'$  for  $P'$  such that  $M' \supseteq M$ .

Two concrete semantics are possible for choice programs: one is an all-answers semantics, and the other is the semantics under which any answer will do – don't care nondeterminism. While an all-answers semantics for choice is not without interesting applications (Greco and Saccà, 1997), the single-answer semantics was adopted by  $\mathcal{L}\mathcal{D}\mathcal{L}++$ , because this is effective at supporting *DB-PTime* problems (Abiteboul *et al.*, 1995). Then, we see that Theorem 2 allows us to compute results incrementally as it is done in differential fixpoint computations; in fact, to find an answer, a program with choice can be implemented as an ordinary program, where the choice predicates are memorized in a table; then newly derived atoms that violate the choice FDs are simply discarded, much in the same way as duplicate atoms are discarded during a fixpoint computation.

Thus positive choice programs represent a class of logic programs that are very well-behaved from both a semantic and a computational viewpoint. The same can be said for choice programs with stratified negation that are defined next.

*Definition 2.2*

Let  $P$  be a choice program with negated goals. Then,  $P$  is said to be stratified when the program obtained from  $P$  by removing its choice goals is stratified.

The stable models for a stratified choice program  $P$  can be computed using an *iterated choice fixpoint* procedure that directly extends the iterated fixpoint procedure for programs with stratified negation (Przymusiński, 1988; Zaniolo et al., 1997); this is summarized next. Let  $P_i$  denote the rules of  $P$  (whose head is) in stratum  $i$ , and let  $P_i^*$  be the union of  $P_j$ ,  $j \leq i$ . Now, if  $M_i$  is a stable model for  $P_i^*$ , then every stable model for  $M_i \cup P_{i+1}$  is a stable model for the program  $P_{i+1}^*$ . Therefore, the stable models of stratified choice programs can be computed by modifying the iterated fixpoint procedure used for stratified programs so that choice models (rather than the least models) are computed for strata containing choice rules (Giannotti et al., 1998).

*The power of choice*

The expressive power of choice was studied in Giannotti et al. (2001b), where it was shown that stratified Datalog with choice can express all computations that are polynomial in the size of the database (i.e. DB-PTIME queries (Abiteboul et al., 1995)). Without choice, DB-PTIME cannot be expressed in stratified Datalog, unless a predefined total order is assumed for the universe, an assumption that would violate the *genericity* principle (Abiteboul et al., 1995). In terms of computational power, non-determinism and order fulfill a similar function (Abiteboul et al., 1995); in fact, the application of choice can also be viewed as non-deterministically and incrementally generating a possible order on the universe – an order that is made explicit by the predicate chain discussed in Example 2.4.

Before moving to Example 2.4, however, we would like to observe that the version of choice supported in  $\mathcal{L}\mathcal{D}\mathcal{L}++$  is more powerful than other nondeterministic constructs, such as the witness operator (Abiteboul et al., 1995), and an earlier version of choice proposed in Krishnamurthy and Naqvi (1998) (called static choice in Giannotti et al. (2001b)). For instance, the following query cannot be expressed in standard Datalog (since the query is nondeterministic) nor it can be expressed by the early version of choice (Krishnamurthy and Naqvi, 1998) or by the witness construct (Abiteboul et al., 1995). These early constructs express nondeterminism in nonrecursive programs, but suffer from inadequate expressive power in recursive programs (Giannotti et al., 2001b). In particular, they cannot express the query in Example 2.3.

*Example 2.3*

*Rooted spanning tree.* We are given an undirected graph where an edge joining two nodes, say  $x$  and  $y$ , is represented by the pair  $g(x, y)$  and  $g(y, x)$ . Then, a spanning tree in this graph, starting from the source node  $a$ , can be constructed by the following program:

```

st(root, a).
st(X, Y) ←  st(., X), g(X, Y), Y ≠ a, Y ≠ X,
             choice((Y), (X)).

```

To illustrate the presence of multiple total choice models for this program, take a simple graph consisting of the following arcs:

```

g(a, b).  g(b, a).
g(b, c).  g(c, b).
g(a, c).  g(c, a).

```

After the exit rule adds  $st(\text{root}, a)$ , the recursive rule could add  $st(a, b)$  and  $st(a, c)$ , along with the two tuples  $\text{chosen}(a, b)$  and  $\text{chosen}(a, c)$  in the *chosen* table. No further arc can be added after those, since the addition of  $st(b, c)$  or  $st(c, b)$  would violate the FD that follows from  $\text{choice}((Y), (X))$  enforced through the *chosen* table. However, since  $st(\text{root}, a)$  was produced by the first rule (the exit rule), rather than the second rule (the recursive choice rule), the table *chosen* contains no tuple with second argument equal to the source node  $a$ . Therefore, to avoid the addition of  $st(c, a)$  or  $st(b, a)$ , the goal  $Y \neq a$  was added to the recursive rule.

By examining all possible solutions, we conclude that this program has three different choice models, for which we list only the *st*-atoms, below:

1.  $st(a, b)$ ,  $st(b, c)$ .
2.  $st(a, b)$ ,  $st(a, c)$ .
3.  $st(a, c)$ ,  $st(c, b)$ .

In addition to supporting *nondeterministic* queries, the introduction of the choice extends the power of Datalog for *deterministic* queries. This can be illustrated by the following choice program that places the elements of a relation  $d(Y)$  into a chain, thus establishing a random total order on these elements; then checks if the last element in the chain is even.

*Example 2.4*

The odd parity query by arranging the elements of a set in a chain. The elements of the set are stored by means of facts of the form  $d(Y)$ .

```

chain(nil, nil).
chain(X, Y) ←  chain(., X), d(Y),
               choice((X), (Y)), choice((Y), (X)).

odd(X) ←      chain(nil, X).
odd(Z) ←      odd(X), chain(X, Y), chain(Y, Z).
isodd ←       odd(X), ¬chain(X, Y).

```

Here  $\text{chain}(\text{nil}, \text{nil})$  is the root of a chain linking all the elements of  $d(Y)$  – thus inducing a total order on elements of  $d$ .

The negated goal in the last rule defines the last element in the chain. Observe that the final *isodd* answer does not depend on the particular chain constructed; it only

depends upon its length that is equal to the cardinality of the set. Thus stratified Datalog with choice can express deterministic queries, such as the parity query, that cannot be expressed in stratified Datalog without choice (Abiteboul *et al.*, 1995).

The parity query cannot be expressed in Datalog with stratified negation unless we assume that the underlying universe is totally ordered – an assumption that violates the data independence principle of *genericity* (Abiteboul *et al.*, 1995). The benefits of this added expressive power in real-life applications follows from the fact that the chain program used in Example 2.4, above, to compute the odd parity query can be used to stream through the elements of a set one by one, and compute arbitrary aggregates on them. For instance, to count the cardinality of the set  $d(Y)$  we can write:

$$\begin{aligned} \text{mcount}(X, 1) &\leftarrow \text{chain}(\text{nil}, X). \\ \text{mcount}(Y, J1) &\leftarrow \text{mcount}(X, J), \text{chain}(X, Y), J1 = J + 1. \\ \text{count}(J) &\leftarrow \text{mcount}(X, J), \neg \text{chain}(X, Y). \end{aligned}$$

The negated goal in the last rule qualifies the element(s)  $X$  without a successor in the chain, i.e.  $X$  for which  $\neg \text{chain}(X, Y)$  holds for all  $Y$ s. Therefore, `count` is defined by a program containing (and stratified with respect to) negation; thus, if `count` is then used as a builtin aggregate, the stratification requirement must be enforced upon every program that uses `count`.

However, if we seek to determine if the base relation  $d(Y)$  has more than 14 elements, then we can use the `mcount` aggregate instead of `count`, as follows:

$$\text{morethan14} \leftarrow \text{mcount}(\_, J), J > 14.$$

Now, `mcount` is what is commonly known as an *online* aggregate (Hellerstein *et al.*, 1997): i.e. an aggregate that produces *early returns* rather than final returns as traditional aggregates. The use of `mcount` over `count` offers clear performance benefits; in fact, the computation of `morethan14` can be terminated after 14 items, whereas the application of `count` requires visiting all the items in the chain. From a logical viewpoint, the benefits are even greater, since `count` is no longer needed and the rule defining it can be eliminated – leaving us with the program defining `mcount`, which is free of negation. Thus, no restriction is needed when using `mcount` in recursive programs; and indeed, `mcount` (and `morethan14`) define monotonic mappings in the lattice of set-containment.

In summary, the use of choice led us to (i) a simple and general definition of the concept of aggregates, including user defined aggregates (UDAs), and (ii) the identification of a special subclass of UDAs that are free from the yoke of stratification, because they are monotonic. This topic is further discussed in the next section.

## 2.2 User defined aggregates

The importance of aggregates in deductive databases has been recognized for a long time (Ross and Sagiv, 1997; Van Gelder, 1993; Kemp and Ramamohanarao, 1998).



In particular, there have been several attempts to overcome the limitations placed on the use of aggregates in programs because of their nonmonotonic nature (Finkelstein, 1996). Of particular interest is the work presented in Ross and Sagiv (1997), where it is shown that rules with aggregates often define monotonic mappings in special lattices, i.e. in lattices different from the standard set-containment lattice used for  $T_p$ . Furthermore, programs with such monotonic aggregates can express many interesting applications (Ross and Sagiv, 1997). Unfortunately, the lattice that makes the aggregate rules of a given program monotonic is very difficult to identify automatically (Van Gelder, 1993); this problem prevents the deployment of such a notion of monotonicity in real deductive database systems.

A new wave of decision support applications has recently underscored the importance of aggregates and the need for a wide range of new aggregates (Han and Kamber, 2001). Examples include rollups and datacubes for OLAP applications, running aggregates and window aggregates in time series analysis, and special versions of standard aggregates used to construct classifiers or association rules in data-mining (Han and Kamber, 2001). Furthermore, a new form of aggregation, called online aggregation, finds many uses in data-intensive applications (Hellerstein *et al.*, 1997). To better serve this wide new assortment of applications requiring specialized aggregates, a deductive database system should support User Defined Aggregates (UDAs). Therefore, the new  $\mathcal{L}\mathcal{D}\mathcal{L}++$  system supports powerful UDAs, including online aggregates and monotonic aggregates, in a simple rule-based framework built on formal logic-based semantics.

In  $\mathcal{L}\mathcal{D}\mathcal{L}++$  users can define a new aggregate by writing the `single`, `multi`, and `freturn` rules (however, `ereturn` rules can be used to supplement or replace `freturn` rules). The `single` rule defines the computation for the first element of the set (for instance `mcount` has its second argument set to 1), while `multi` defines the induction step whereby the value of the aggregate on a set of  $n + 1$  elements is derived from the aggregate value of the previous set with  $n$  elements and the value of  $(n + 1)$ th element itself. A unique aggregate name is used as the first argument in the head of these rules to eliminate any interference between the rules defining different aggregates. For instance, for computing averages we must compute both the count and the sum of the elements seen so far:

```
single(avg, Y, cs(1, Y)).
multi(avg, Y, cs(Cnt, Sum), cs(Cnt1, Sum1)) ←
    Cnt1 = Cnt + 1, Sum1 = Sum + Y.
```

Then, we write a `freturn` rule that upon visiting the final element in  $d(Y)$  produces the ratio of sum over count, as follows:

```
freturn(avg, Y, cs(Cnt, Sum), Val) ← Val = Sum/Cnt.
```

After an aggregate is defined by its `single`, `multi`, `ereturn` and/or `freturn` rules, it can be invoked and used in other rules. For instance, our newly defined `avg` can be invoked as follows:

```
p(avg(Y)) ← d(Y).
```

Thus  $\mathcal{L}\mathcal{D}\mathcal{L}++$  uses the special notation of pointed brackets, in the head of rules, to denote the application of an aggregate. This syntax, that has been adopted by other languages (Ramakrishnan *et al.*, 1993), also supports an implicit ‘group by’ construct, whereby the aggregate arguments in the head are implicitly grouped by the other arguments in the head. Thus, to find the average salary of employees grouped by department a user can write the following rule:

$$\text{davg}(\text{DeptNo}, \text{avg}(\text{Sal})) \leftarrow \text{employee}(\text{Eno}, \text{Sal}, \text{DeptNo}).$$

The formal semantics of UDAs was introduced in Zaniolo and Wang (1999) and is described in the Appendix: basically, the aggregate invocation rules and the aggregate definition rules are rewritten into an equivalent program that calls on the chain predicate defined as in Example 2.4. (Naturally, for the sake of efficiency, the  $\mathcal{L}\mathcal{D}\mathcal{L}++$  system shortcuts the full rewriting used to define their formal semantics, and implement the UDAs by a more direct implementation.)

$\mathcal{L}\mathcal{D}\mathcal{L}++$  UDAs have also been extended to support online aggregation (Hellerstein *et al.*, 1997). This is achieved by using `ereturn` rules in the definition of UDAs, to either supplement, or replace `freturn` rules.

For example, the computation of averages normally produces an approximate value long before the whole data set is visited. Then, we might want to see the average value obtained so far every 100 elements. Then, the following rule will be added:

$$\begin{aligned} \text{ereturn}(\text{avg}, X, (\text{Sum}, \text{Count}), \text{Avg}) \leftarrow \\ \text{Count mod } 100 = 0, \text{Avg} = \text{Sum}/\text{Count}. \end{aligned}$$

Thus the `ereturn` rules produce *early returns*, while the `freturn` rules produce final returns.

As second example, let us consider the well-known problem of coalescing after temporal projection in temporal databases (Zaniolo *et al.*, 1997). For instance in Example 2.5, after projecting out from the employee relation the salary column, we might have a situation where the same `Eno` appears in tuples where their valid-time intervals overlap; then these intervals must be coalesced. Here, we use closed intervals represented by the pair `(From, To)` where `From` is the start-time, and `To` is the end-time. Under the assumption that tuples are sorted by increasing start-time, then we can use a special `coales` aggregate to perform the task in one pass through the data.

#### Example 2.5

Coalescing overlapping intervals sorted by start time.

$$\begin{aligned} \text{empProj}(\text{Eno}, \text{coales}(\langle \text{From}, \text{To} \rangle)) \leftarrow \text{emp}(\text{Eno}, \rightarrow, \rightarrow, (\text{From}, \text{To})). \\ \text{single}(\text{coales}, (\text{Frm}, \text{To}), (\text{Frm}, \text{To})). \\ \text{multi}(\text{coales}, (\text{Nfr}, \text{Nto}), (\text{Cfr}, \text{Cto}), (\text{Cfr}, \text{Nto})) \leftarrow \\ \text{Nfr} \leq \text{Cto}, \text{Nto} > \text{Cto}. \end{aligned}$$

```

multi(coales, (Nfr, Nto), (Cfr, Cto), (Cfr, Cto)) ←
    Nfr ≤ Cto, Nto ≤ Cto.
multi(coales, (Nfr, Nto), (Cfr, Cto), (Nfr, Nto)) ← Cto < Nfr.

ereturn(coales, (Nfr, Nto), (Cfr, Cto), (Cfr, Cto)) ← Cto < Nfr.
freturn(coales, _, LastInt, LastInt).

```

Since the input intervals are ordered by their start time, the new interval  $(Nfr, Nto)$  overlaps the current interval  $(Cfr, Cto)$  when  $Nfr \leq Cto$ ; in this situation, the two intervals are merged into one that begins at  $Cfr$  and ends with the larger of  $Nto$  and  $Cto$ . When, the new interval does not overlap with the current interval, this is returned by the `ereturn` rule, while the new interval becomes the current one (see the last `multi` rule).

Let  $P$  be a program. A rule  $r$  of  $P$  whose head contains aggregates is called an *aggregate rule*. Then,  $P$  is said to be stratified w.r.t. aggregates when for each aggregate rule  $r$  in  $P$ , the stratum of  $r$ 's head predicate is strictly higher than the stratum of each predicate in the head of  $r$ . Therefore, the previous program is stratified with respect to `coales` which is nonmonotonic since it uses both early returns and final returns.

While, programs stratified with respect to aggregates can be used in many applications, more advanced applications require the use of aggregates in more general settings. Thus,  $\mathcal{L}\mathcal{D}\mathcal{L}++$  supports the usage of arbitrary aggregates in XY-stratified programs, which will be discussed in Section 3. Furthermore  $\mathcal{L}\mathcal{D}\mathcal{L}++$  supports the monotonic aggregates that can be used freely in recursion.

#### Monotone aggregation

An important result that follows from the formalization of the semantics of UDAs (Zaniolo and Wang, 1999) (see also the Appendix), is that UDA defined without final return rules, i.e. no `freturn` rule, define monotonic mappings, and can thus be used without restrictions in the definition of recursive predicates. For instance, we will next define a continuous count that returns the current count after each new element (thus final returns are here omitted since they are redundant).

```

single(mcount, Y, 1).
multi(mcount, Y, Old, New) ← New = Old + 1.
ereturn(mcount, Y, Old, New) ← New = Old + 1.

```

Monotonic aggregates allow us to express the following two examples taken from Ross and Sagiv (1997).

*Join the Party.* Some people will come to the party no matter what, and their names are stored in a `sure(Person)` relation. But others will join only after they know that at least  $K = 3$  of their friends will be there. Here, `friend(P, F)` denotes that  $F$  is a friend of person  $P$ .

```

willcome(P) ← sure(P).
willcome(P) ← c_friends(P, K), K ≥ 3.
c_friends(P, mcount(F)) ← willcome(F), friend(P, F).

```

Consider now a computation of these rules on the following database.

```

friend(jerry,mark).      sure(mark).
friend(penny,mark).     sure(tom).
friend(jerry,jane).     sure(jane).
friend(penny,jane).
friend(jerry,penny).
friend(penny,tom).

```

Then, the basic semi-naive computation yields:

```

willcome(mark),willcome(tom),willcome(jane),
c_friends(jerry,1),c_friends(penny,1),c_friends(jerry,2),
c_friends(penny,2),c_friends(penny,3),willcome(penny),
c_friends(jerry,3),willcome(jerry).

```

This example illustrates how the standard semi-naive computation can be applied to queries containing monotone UDAs. Another interesting example is transitive ownership and control of corporations.

*Company Control.* Say that  $\text{owns}(C1,C2,Per)$  denotes the percentage of shares that corporation  $C1$  owns of corporation  $C2$ . Then,  $C1$  controls  $C2$  if it owns more than, say, 50% of its shares. In general, to decide whether  $C1$  controls  $C3$  we must also add the shares owned by corporations, such as  $C2$ , that are controlled by  $C1$ . This yields the transitive control rules defined with the help of a continuous sum aggregate that returns the partial sum for each new element:

```

control(C,C) ←      owns(C,-,-).
control(Onr,C) ←   towns(Onr,C,Per),Per > 50.
towns(Onr,C2,msum(Per)) ← control(Onr,C1),owns(C1,C2,Per).

single(msum,Y,Y).
multi(msum,Y,Old,New) ←   New = Old + Y.
ereturn(msum,Y,Old,New) ← New = Old + Y.

```

Thus, every company controls itself, and a company  $C1$  that has transitive ownership of more than 50% of  $C2$ 's shares controls  $C2$ . In the last rule,  $\text{towns}$  computes transitive ownership with the help of  $\text{msum}$  that adds up the shares of controlling companies. Observe that any pair  $(\text{Onr},C2)$  is added at most once to  $\text{control}$ , thus the contribution of  $C1$  to  $\text{Onr}$ 's transitive ownership of  $C2$  is only accounted once.

*Bill-of-Materials (BoM) Applications.* BoM applications represent an important application area that requires aggregates in recursive rules. For instance, let us say that  $\text{assembly}(P1,P2,QT)$  denotes that  $P1$  contains part  $P2$  in quantity  $QT$ . We also have elementary parts described by the relation  $\text{basic\_part}(\text{Part},\text{Price})$ . Then, the

following program computes the cost of a part as the sum of the cost of the basic parts it contains:

```

part_cost(Part, 0, Cst) ←      basic_part(Part, Cst).
part_cost(Part, mcount(Sb), msum(MCst)) ←
    part_cost(Sb, ChC, Cst), prolfc(Sb, ChC),
    assembly(Part, Sb, Mult), MCst = Cst * Mult.

```

Thus, the key condition in the body of the second rule is that a subpart  $Sb$  is counted in `part_cost` only when all  $Sb$ 's children have been counted. This occurs when the number of  $Sb$ 's children counted so far by `mcount` is equal to the out-degree of this node in the graph representing assembly. This number is kept in the prolificacy table, `prolfc(Part, ChC)`, which can be computed as follows:

```

prolfc(P1, count(P2)) ←  assembly(P1, P2, _).
prolfc(P1, 0) ←          basic_part(P1, _).

```

Therefore the simple and general solution of the monotonic aggregation problem introduced by  $\mathcal{L}\mathcal{D}\mathcal{L}++$  allows the concise expression of many interesting algorithms. This concept can also be extended easily to SQL recursive queries, as discussed in (Wang and Zaniolo, 2000) where additional applications are also discussed.

### 2.3 Beyond stratification

The need to go beyond stratification has motivated much recent research. Several deductive database systems have addressed it by supporting the notion of modular stratification (Ross, 1994). Unfortunately, this approach suffers from poor usability, since the existence of a modular stratification for a program can depend on its extensional information (i.e. its fact base) and, in general, cannot be checked without executing the program. The standard notion of stratification is instead much easier to use, since it provides a simple criterion for the programmer to follow and for the compiler to use when validating the program and optimizing its execution. Therefore,  $\mathcal{L}\mathcal{D}\mathcal{L}++$  has introduced the notion of XY-stratified programs that preserves the compilability and usability benefits of stratified programs while achieving the expressive power of well-founded models (Kemp *et al.*, 1995). XY-stratified programs are locally stratified explicitly by a temporal argument: thus, they can be viewed as  $\text{Datalog}_{1S}$  programs, which are known to provide a powerful tool for temporal reasoning (Baudinet *et al.*, 1994; Zaniolo *et al.*, 1997), or as Statelog programs that were used to model active databases (Lausen *et al.*, 1998b). The deductive database system Aditi (Kemp and Ramamohanarao, 1998) also supports the closely related concept of explicitly locally stratified programs, which were shown to be as powerful as well-founded models, since they can express their alternating fixpoint computation (Kemp *et al.*, 1995).

For instance, the ancestors of `marc`, with the number of generations that separate them from `marc`, can be computed using the following program which models the differential fixpoint computation:

*Example 2.6*

Computing ancestors of Marc and their remoteness from Marc using differential fixpoint approach.

```

r1 : delta_anc(0,marc).
r2 : delta_anc(J + 1, Y) ← delta_anc(J, X), parent(Y, X),
                             ¬all_anc(J, Y).
r3 : all_anc(J + 1, X) ← all_anc(J, X).
r4 : all_anc(J, X) ← delta_anc(J, X).

```

This program is locally stratified by the first arguments in `delta_anc` and `all_anc` that serve as temporal arguments (thus `+1` is a postfix successor function symbol, much the same as  $s(J)$  that denotes the successor of  $J$  in  $\text{Datalog}_{1S}$  (Zaniolo et al., 1997)). The zero stratum consists of atoms of nonrecursive predicates such as `parent` and of atoms that unify with `all_anc(0, X)` or `delta_anc(0, X)`. The  $k^{\text{th}}$  stratum consists of atoms of the form `all_anc(k, X)`, `delta_anc(k, X)`. Thus, the previous program is locally stratified (Przymusiński, 1988), since the heads of recursive rules belong to strata that are one above those of their goals. Alternatively, we can view the previous program as a compact representation for the stratified program obtained by instantiating the temporal argument to integers and attaching them to the predicate names, thus generating an infinite sequence of unique names.

Also, observe that the temporal arguments in rules are either the same as, or one less than, the temporal argument in the head. Then, there are two kinds of rules in our example: (i) X-rules (i.e. a horizontal rules) where the temporal argument in each of their goals is the same as that in their heads, and (ii) Y-rules (i.e. a vertical rules) where the temporal arguments in some of their goals are one less than those in their heads. Formally, let  $P$  be a set of rules defining mutually recursive predicates, where each recursive predicate has a distinguished temporal argument and every rule in  $P$  is either an X-rule or a Y-rule. Then,  $P$  will be said to be an XY-program. For instance, the program in Example 2.6 is an XY-program, where  $r_4$  and  $r_1$  are X-rules, while  $r_2$  and  $r_3$  are Y-rules.

A simple test can now be used to decide whether an XY-program  $P$  is locally stratified. The test begins by labelling all the head predicates in  $P$  with the prefix ‘new’. Then, the body predicates with the same temporal argument as the head are also labelled with the prefix ‘new’, while the others are labelled with the prefix ‘old’. Finally, the temporal arguments are dropped from the program. The resulting program is called the *bistate version* of  $P$  and is denoted  $P_{bis}$ .

*Example 2.7*

The bistate version of the program in Example 2.6:

```

new_delta_anc(marc).
new_delta_anc(Y) ← old_delta_anc(X), parent(Y, X),
                   ¬old_all_anc(Y).
new_all_anc(X) ← new_delta_anc(X).
new_all_anc(X) ← old_all_anc(X).

```

Now we have that (Zaniolo et al., 1993):

*Definition 2.3*

Let  $P$  be an XY-program.  $P$  is said to be XY-stratified when  $P_{bis}$  is a stratified program.

*Theorem 2.3*

Let  $P$  be an XY-stratified program. Then  $P$  is locally stratified.

The program of Example 2.7 is stratified with the following strata:  $S_0 = \{\text{parent}, \text{old\_all\_anc}, \text{old\_delta\_anc}\}$ ,  $S_1 = \{\text{new\_delta\_anc}\}$ , and  $S_2 = \{\text{new\_all\_anc}\}$ . Thus, the program in Example 2.6 is locally stratified.

For an XY-stratified program  $P$ , the general iterated fixpoint procedure (Przymusiński, 1988) used to compute the stable model of locally stratified programs (Zaniolo *et al.*, 1993) becomes quite simple; basically it reduces to a repeated computation over the stratified program  $P_{bis}$ . For instance, for Example 2.7 we compute `new_delta_anc` from `old_delta_anc` and then `new_all_anc` from this. Then, the ‘old’ relations are re-initialized with the content of the ‘new’ ones so derived, and the process is repeated. Furthermore, since the temporal arguments have been removed from this program, we need to

1. store the temporal argument as an external fact `counter(T)`,
2. add a new goal `counter(Ir)` to each exit rule  $r$  in  $P_{bis}$ , where  $I_r$  is the variable from the temporal arguments of the original rule  $r$ , and
3. For each recursive predicate  $q$  add the rule:

$$q(J, X) \leftarrow \text{new\_}q(X), \text{counter}(J).$$

The program so constructed will be called the *synchronized* bistate version of  $P$ , denoted  $\text{syncbi}(P)$ . For instance, to obtain the synchronized version of the program in Example 2.7, we need to change the first rule to

$$\text{new\_delta\_anc}(\text{marc}) \leftarrow \text{counter}(0).$$

since the temporal argument in the original exit rule was the constant 0. Then, we must add the following rules:

$$\begin{aligned} \text{delta\_anc}(J, X) &\leftarrow \text{new\_delta\_anc}(X), \text{counter}(J). \\ \text{all\_anc}(J, X) &\leftarrow \text{new\_all\_anc}(X), \text{counter}(J). \end{aligned}$$

Then, the iterated fixpoint computation for an XY-stratified program can be implemented by the following procedure:

*Procedure 2.1*

*Computing a stable model of an XY-stratified program  $P$* : Add the fact `counter(0)`. Then, forever repeat the following two steps:

1. Compute the stable model of  $\text{syncbi}(P)$ .
2. For each recursive predicate  $q$ , replace `old_q` with `new_q`, computed in the previous step. Then, increase the value of `counter` by one.

Since  $\text{syncbi}(P)$  is stratified, we can then use the iterated fixpoint computation to compute its stable model.

Since each XY-stratified program is locally stratified (Przymusiński, 1988), it is guaranteed to have a unique stable model, which is also known as its perfect model (Przymusiński, 1988). However, the special syntactic structure of XY-stratified programs allows an efficient computation of their perfect models using Procedure 4; moreover, in the actual  $\mathcal{L}\mathcal{D}\mathcal{L}++$  implementation, this computation is further improved with the optimization techniques discussed next. For instance, the replacement of `old_q` with `new_q` described in the last step of Procedure 2.1 becomes an operation of (small) constant cost when it is implemented by switching the pointers to the relations. A second improvement concerns *copy rules*, such as the last rule in Example 2.6. For instance  $r_3$  in Example 2.6 is a copy rule that copies the new values of `all_anc` from its old values. Observe that the body and the head of this rule are identical, except for the prefixes `new` or `old`, in its bistate version (Example 2.7). Thus, in order to compute `new_all_anc`, we first execute the copy rule by simply setting the pointer to `new_all_anc` to point to `old_all_anc` – a constant-time operation. Rule  $r_4$  that adds tuples to `new_all_anc` is then executed after  $r_3$ .

In writing XY-stratified programs, the user must also be concerned with termination conditions since, for example, a rule such as  $r_3$  in Example 2.6 could, if left unchecked, keep producing `all_anc` results under a new temporal argument, after `delta` becomes empty. One solution to this problem is for the user to add the goal `delta_anc(J, _)` to rule  $r_3$ . Then, the computation `all_anc` stops as soon as no new `delta_anc(J, _)` is generated. Alternatively, our program could be called from a goal such as `delta_anc(J, Y)`. In this case, if  $r_2$  fails to produce any result for a value `J`, no more results can be produced at successive steps, since `delta_anc(J, Y)` is a positive goal of  $r_2$ . The  $\mathcal{L}\mathcal{D}\mathcal{L}++$  system is capable of recognizing these situations, and it will terminate the computation of Procedure 2.1 when either condition occurs.

Example 2.8 solves the coalescing problem without relying on tuples being sorted on their start-time – an assumption made in Example 2.5. Therefore, we use the XY-stratified program of Example 2.8, which iterates over two basic computation steps. The first step is defined by the `overlap` rule, which identifies pairs of distinct intervals that overlap, where the first interval contains the start of the second interval. The second step consists of deriving a new interval that begins at the start of the first interval, and ends at the later of the two endpoints. Finally, a rule `final_e_hist` returns the intervals that do not overlap other intervals (after eliminating the temporal argument).

#### Example 2.8

Coalescing overlapping periods into maximal periods after a projection:

```
e_hist(0, Eno, Frm, To) ← emp_dep_sal(0, Eno, _, _, Frm, To).
overlap(J + 1, Eno, Frm1, To1, Frm2, To2) ←
    e_hist(J, Eno, Frm1, To1),
    e_hist(J, Eno, Frm2, To2),
    Frm1 ≤ Frm2, Frm2 ≤ To1,
    distinct(Frm1, To1, Frm2, To2).
e_hist(J, Eno, Frm1, To) ← overlap(J, Eno, Frm1, To1, Frm2, To2),
    select_larger(To1, To2, To).
```



```

final_e_hist(J + 1, Eno, Frm, To) ←  e_hist(J, Eno, Frm, To),
                                     ¬overlap(J + 1, Eno, Frm, To, →, →).

distinct(Frm1, To1, Frm2, To2) ←  To1 ≠ To2.
distinct(Frm1, To1, Frm2, To2) ←  Frm1 ≠ Frm2.
select_larger(X, Y, X) ←  X ≥ Y.
select_larger(X, Y, Y) ←  Y > X.

```

As demonstrated by these examples, XY-stratified programs allow an efficient logic-based expression of procedural algorithms. For instance, the alternating fixpoint procedure used in the computation of well-founded models can also be expressed using these programs (Kemp *et al.*, 1995). In general, XY-stratified programs are quite powerful, as demonstrated by fact that these programs (without choice, aggregates, and function symbol) are known to be equivalent to Statelog programs (Lausen *et al.*, 1998a), which have PSPACE complexity and can express the WHILE queries (Abiteboul *et al.*, 1995). Finally, observe that the bistrate programs for the examples used here are nonrecursive. In general, by making the computation of the recursive predicate explicit as it was done for the anc example, it is possible to rewrite an XY-stratified program  $P$  whose bistrate version  $P_{bis}$  is recursive into an XY-stratified program  $P'$  whose bistrate version  $P'_{bis}$  is nonrecursive.

#### Choice and aggregates in XY-stratified programs

As described in section 2.1, choice can be used in stratified programs with no restriction, and its stable model can be computed by an iterated choice fixpoint procedure. Generalizing such notion, the  $\mathcal{L}\mathcal{D}\mathcal{L}++$  system supports the use of choice in programs that are XY-stratified with respect to negation. The following conditions are, however, enforced to assure the existence of stable models for a given program  $P$  (Giannotti *et al.*, 1998):

- the program obtained from  $P$  by removing its choice goals is XY-stratified w.r.t. negation, and
- if  $r$  is a recursive choice rule in  $P$ , then *some* choice goal of  $r$  contains  $r$ 's temporal variable in its left side.

After checking these conditions, the  $\mathcal{L}\mathcal{D}\mathcal{L}++$  compiler constructs  $synchi(P)$  by dropping the temporal variable from the choice goals and transforming the rest of the rules as described in the previous section. Then, the program  $synchi(P)$  so obtained is a stratified choice program and its stable models can be computed accordingly; therefore, each stable model for the original XY-stratified program  $P$  is computed by simply applying Procedure 2.1 with no modification (Zaniolo *et al.*, 1997; Giannotti *et al.*, 1998).

Using the simple syntactic characterization given in section 2.2,  $\mathcal{L}\mathcal{D}\mathcal{L}++$  draws a sharp distinction between monotonic and nonmonotonic aggregates. No restriction is imposed on programs with only monotonic aggregates and no negation. But recursive programs with nonmonotonic aggregates must satisfy the following conditions (which assure that once the aggregates are expanded as described in

section 2.2 the resulting choice program satisfies the XY-stratification conditions for choice programs discussed in the previous paragraph):

- For each recursive rule, the temporal variable must be contained in the group-by attributes.
- The bistrate version of  $P$  must be stratified w.r.t. negation and nonmonotonic aggregates, and

After checking these simple conditions, the  $\mathcal{L}\mathcal{D}\mathcal{L}++$  compiler proceeds with the usual computation of  $\text{syncbi}(P)$  as previously described.

For instance, the following XY-stratified program with aggregates expresses Floyd's algorithm to compute the least-cost path between pairs of nodes in a graph. Here,  $g(X, Y, C)$  denotes an arc from  $X$  to  $Y$  of cost  $C$ :

*Example 2.9*

Floyd's least-cost paths between all node pairs.

$$\begin{aligned} \text{delta}(0, X, Y, C) &\leftarrow g(X, Y, C). \\ \text{newmin}(J, X, Z, \min(C)) &\leftarrow \text{new}(J, X, Z, C). \\ \text{discard}(J, X, Z, C) &\leftarrow \text{newmin}(J, X, Z, C1), \text{all}(J, X, Z, C2), C1 \geq C2. \\ \text{delta}(J, X, Z, C) &\leftarrow \text{newmin}(J, X, X, C), \neg \text{discard}(J, X, Z, ). \\ \text{all}(J, X, Z, C) &\leftarrow \text{delta}(J, X, Z, C). \end{aligned}$$

The fourth rule in this example uses a nonmonotonic min aggregate to select the least cost pairs among those just generated (observe that the temporal variable  $J$  appears among the group-by attributes). The next two rules derive the new delta pairs by discarding from new those that are larger than any existing pair in all. This new delta is then used to update all and compute new pairs.

By supporting UDAs, choice, and XY-stratification  $\mathcal{L}\mathcal{D}\mathcal{L}++$  provides a powerful, fully integrated framework for expressing logic-based computation and modelling. In addition to express complex computations (Zaniolo et al., 1998), this power has been used to model the AI planning problem (Brogi et al., 1997), database updates, and active database rules (Zaniolo, 1997). For instance, to model AI planning, preconditions can simply be expressed by rules, choice can be used to select among applicable actions, and frame axioms can be expressed by XY-stratified rules that describe changes from the old state to the new state (Brogi et al., 1997).

### 3 The system

The main objectives in the design of the  $\mathcal{L}\mathcal{D}\mathcal{L}++$  system, were (i) strengthening the architecture of the previous  $\mathcal{L}\mathcal{D}\mathcal{L}$  system (Chimenti et al., 1990), (ii) improving the system's usability and the application development turnaround time, and (iii) provide efficient support for the new language constructs.

While the first objective could be achieved by building on and extending the general architecture of the predecessor  $\mathcal{L}\mathcal{D}\mathcal{L}$  system, the second objective forced us to depart significantly from the compilation and execution approach used by the  $\mathcal{L}\mathcal{D}\mathcal{L}$  system. In fact, the old system adhered closely to the set-oriented

semantics of relational algebra and relational databases; therefore, it computed and accumulated all partial results before returning the whole set to the user. However, our experience in developing applications indicated that a more interactive and incremental computation model was preferable, i.e. one where users see the results incrementally as they are produced. This allows developers to monitor better the computation as it progresses, helping them debugging their programs, and, e.g. allowing them to stop promptly executions that have fallen into infinite loops.

Therefore,  $\mathcal{L}\mathcal{D}\mathcal{L}++$  uses a pipelined execution model, whereby tuples are generated one at a time as they are needed (i.e. lazily as the consumer requests them, rather than eagerly). This approach also realizes objective (iii) by providing better support for new constructs, such as choice and on-line aggregation, and for intelligent backtracking optimization (discussed in the next section).

The  $\mathcal{L}\mathcal{D}\mathcal{L}++$  system also adopted a shallow-compilation approach that achieves faster turnaround during program development and enhances the overall usability; this approach also made it easier to support on-line debugging and meta-level extensions. The previous  $\mathcal{L}\mathcal{D}\mathcal{L}$  system was instead optimized for performance; thus, it used a deep-compilation approach where the original program was translated into a (large) C program – whose compilation and linking slowed the development turnaround time. The architecture of the system is summarized in the next section; additional information, a web demo, and instructions on downloading for noncommercial use can be found in Zaniolo (1998).

### 3.1 Architecture

The overall architecture of the  $\mathcal{L}\mathcal{D}\mathcal{L}++$  system and its main components are shown in figure 1. The major components of the system are:

*The compiler.* The compiler reads in  $\mathcal{L}\mathcal{D}\mathcal{L}++$  programs and constructs the *Global Predicate Connection Graph* (PCG). For each query form, the compiler partially evaluates the PCG, transforming it into a network of objects that are executed by the interpreter. The compiler is basically similar to that of the old system (Chimenti *et al.*, 1990), and is responsible for checking the safety of queries, and rewriting the recursive rules using techniques such the Magic Sets method (Bancilhon *et al.*, 1986), and the more specialized methods for left-linear and right-linear rules (Ullman, 1989). These rewriting techniques result in an efficient execution plan for queries.

*The database managers.* The  $\mathcal{L}\mathcal{D}\mathcal{L}$  experience confirmed the desirability supporting access to (i) an internal (fast-path) database and (ii) multiple external DBMSs in a transparent fashion. This led to the design of a new system where the two types of database managers are fully integrated.

The internal database is shown in figure 1 as Fact Base Manager. This module supports the management and retrieval of  $\mathcal{L}\mathcal{D}\mathcal{L}++$  complex objects, including sets and lists, and of temporary relations obtained during the computation. In addition to supporting users' data defined by the schema as internal relations, the

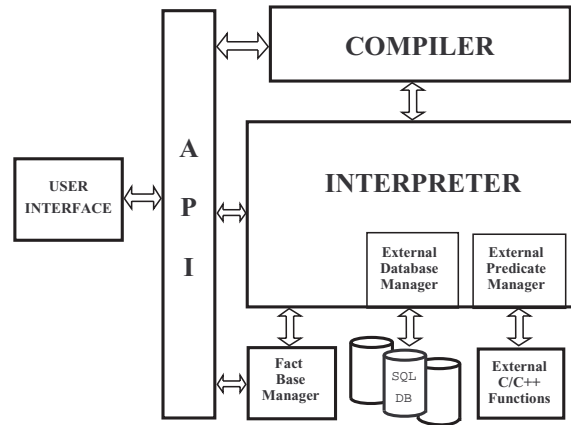


Fig. 1.  $\mathcal{L}\mathcal{D}\mathcal{L}++$  open architecture.

interpreter relies on the local database to store and manage temporary data sets. The internal database is designed as a virtual-memory record manager: thus its internal organization and indexing schemes are optimized for the situation where the pages containing frequently used data can reside in main memory. Data is written back onto disk at the commit point of each update transaction; when the transaction aborts the old data is instead restored from disk.

The system also supports an external database manager, which is designed to optimize access to external SQL databases; this is described in section 3.3.

*Interpreter.* The interpreter receives as input a graph of executable objects corresponding to an  $\mathcal{L}\mathcal{D}\mathcal{L}++$  query form generated by the compiler, and executes it by issuing get-next, and other calls, to the local database. Similar calls are also issued by the External Database Manager and the External Predicate Manager to, respectively, external databases, and external functions or software packages that follow the C/C++ calling conventions. Details on the interpreter are presented in the next section.

*User interface.* All applications written in C/C++ can call the  $\mathcal{L}\mathcal{D}\mathcal{L}++$  system via a standard API; thus applications written in  $\mathcal{L}\mathcal{D}\mathcal{L}++$  can be embedded in other procedural systems.

One such application is a line-oriented command interpreter supporting a set of predefined user commands, command completion and on-line help. The command interpreter is supplied with the system, although it is not part of the core system. Basically, the interface is an application built in C++ that can be replaced with other front-ends, including graphical ones based on a GUI, without requiring any changes to the internals of the system. In particular, a Java-based interface for remote demoing was added recently (Zaniolo et al., 1998).

### 3.2 Execution model and interpreter

The abstract machine for the  $\mathcal{L}\mathcal{D}\mathcal{L}++$  interpreter is based upon the architecture described in (Chimenti *et al.*, 1989). An  $\mathcal{L}\mathcal{D}\mathcal{L}++$  program is transformed into a network of active objects and the graph-based interpreter then processes these objects.

*Code generation and execution.* Given a query form, an  $\mathcal{L}\mathcal{D}\mathcal{L}++$  program is transformed into a Predicate Connection Graph (PCG), which can be viewed as an AND/OR graph with annotations. An OR-node represents a predicate occurrence and each AND node represents the head of a rule. The PCG is subsequently compiled into an evaluable data structure called a LAM (for  $\mathcal{L}\mathcal{D}\mathcal{L}++$  Abstract Machine), whose nodes are implemented as instances of C++ classes. Arguments are passed from one node to the other by means of variables. Unification is done at compile time and the sharing of variables avoids useless assignments.

Each node of the generated LAM structure has a virtual<sup>2</sup> ‘GetTuple’ interface, which evaluates the corresponding predicate in the program. Each node also stores a state variable that determines whether this node is being ‘entered’ or is being ‘backtracked’ into. The implementation of this ‘GetTuple’ interface depends upon the type of node. The most basic C++ classes are OR-nodes and AND-nodes; then there are several more specialized subclasses of these two basic types. Such subclasses include the special OR-node that serves as the distinguished *root* node for the query form, internal relations AND-nodes, external relations AND-nodes, etc.

*And/OR graph.* For a generic OR node corresponding to a derived relation, the ‘GetTuple’ interface merely issues ‘GetTuple’ calls to its children (AND nodes). Each successful invocation automatically instantiates the variables of both the child (AND node) and the parent (OR node). Upon backtracking, the last AND node which was successfully executed is executed again. The ‘GetTuple’ on an OR node fails when its last AND node child fails.

The *Dataflow points* represent different entries into the AND/OR nodes, each entry corresponding to a different state of the computation. The dataflow points associated with each node are shown in Table 1 (observe their similarity to ports in Byrd’s Prolog execution model (Byrd, 1980)).

A dataflow point of a node can be directed to a dataflow point of a different node by a *dataflow destination*. The *entry destination* ( $e\_dest$ ) of a given node is the dataflow point to which its entry point is directed. Similarly, *backtrack* ( $b\_dest$ ), *success* ( $s\_dest$ ), and *fail destinations* ( $f\_dest$ ) can be defined. The dataflow destinations represent logical operations between the nodes involved; for example a join or union of the two nodes. The dataflow points and destinations of a node describe how the tuples of that node are combined with tuples from other nodes (but not how those tuples are generated).

<sup>2</sup> Similar to a C++ virtual function.

Table 1. Dataflow points associated with each node

Dataflow	Point	State of Computation
ENTRY	<i>e_dest</i>	getting first tuple of node
BACKTRACK	<i>b_dest</i>	getting next tuple of node
SUCCESS	<i>s_dest</i>	a tuple has been generated
FAIL	<i>f_dest</i>	no more tuples can be generated

To obtain the first tuple of an OR node we get the first tuple of its first child AND node. To obtain the next tuple from an OR node we request it from the AND node that generated the previous tuple. Observe that the currently ‘active’ AND node must be determined at run-time. When no more tuples can be generated for a given AND node, then we go to the next AND node, till the last child AND node is reached. (At this point no more tuples can be generated for the OR node.) Thus, we have:

**OR nodes:** *e\_dest*: the *e\_dest* of the first child AND-node  
*b\_dest*: the *b\_dest* of the “active” child AND node  
*f\_dest*: **if** node is first OR node in rule  
**then** the *f\_dest* point of parent AND node  
**else** the *b\_dest* of previous OR node  
*s\_dest*: **if** node is last OR node in a rule  
**then** the *s\_dest* of parent AND node  
**else** the *e\_dest* of next OR node.

The execution of an AND node is conceptually less complicated. Intuitively, the execution corresponds to a nested loop, where, for each tuple of the first OR node, we generate all matching tuples from the next OR node. This continues until we reach the last OR node. Thus, when generating the next tuple of an AND node, we generate the next matching tuple from the last OR node. If there are no more matching tuples, we generate the next tuple from the previous OR node. When there are no more tuples to be generated by the first OR node, we can generate no more tuples for the AND node. Thus we have:

**AND nodes:** *e\_dest*: the *e\_dest* of first OR child  
*b\_dest*: the *b\_dest* of last OR child  
*f\_dest*: **if** node is last AND child  
**then** *f\_dest* of parent OR node  
**else** *e\_dest* of next AND node  
*s\_dest*: *s\_dest* of parent OR node.

Given a query, the  $\mathcal{L}\mathcal{D}\mathcal{L}++$  system first finds the appropriate LAM graph for the matching query form, then stores any constant being passed to the query form by initializing the variables attached to the *root* node of the LAM graph. Finally, the system begins the execution by repeatedly calling the ‘GetTuple’ method on the root of this graph. When the call fails the execution is complete.

- Step 1.** Move the cursor  $C_j$  to the next tuple of  $R$ , and consume the tuple.  
**Step 2.** If Step 1 fails (thus,  $C_j$  is the last tuple of  $R$ ), check the fixpoint flag  $F$ .  
**Step 3.** If the fixpoint is reached, return failure.  
**Step 4.** If the fixpoint is not reached, call the current rule to generate a new tuple.  
**Step 5.** If a new tuple is generated, add it to the relation  $R$ , advance  $C_j$  and return the tuple.  
**Step 6.** Otherwise, repeat Step 2.

Fig. 2. Lazy fixpoint producer.

*Lazy Evaluation of Fixpoints.*  $\mathcal{L}\mathcal{D}\mathcal{L}++$  adopts a lazy evaluation approach (*pipelining*) as its primary execution model, which is naturally supported by the AND/OR graph described above. This model is also supported through the lazy evaluation of fixpoints. The traditional implementation of fixpoints (Ullman, 1989; Zaniolo *et al.*, 1997) assumes an eager computation, where new tuples are generated until the fixpoint is reached.  $\mathcal{L}\mathcal{D}\mathcal{L}++$  instead supports lazy computation where the recursive rules produce new tuples only in response to the goal that, as a consumer, calls the recursive predicate. Multiple consumers can be served by one producer, since each consumer  $j$  uses a separate cursor  $C_j$  to access the relation  $R$  written by the producer. Whenever  $j$  needs a new tuple, it proceeds as shown in figure 2.

A limitation of pipelining is that the internal state of each node must be kept for computation to resume where the last call left off. This creates a problem when several goals call the same predicate (i.e. the same subtree in the PCG is shared). Multiple invocations of a shared node can interfere with each other (non-reentrant code). Solutions to this problem include (i) using a stack as in Prolog, and (ii) duplicating the source code as in the  $\mathcal{L}\mathcal{D}\mathcal{L}$  system – thus ensuring that the PCG is a tree, rather than a DAG (Chimenti *et al.*, 1990). In the  $\mathcal{L}\mathcal{D}\mathcal{L}++$  system, we instead use the lazy producer approach described above for situations where the calling goals have no bound argument. If there are bound arguments in consuming predicates we duplicate the node. However, since each node is implemented as a C++ class, we simply generate multiple instances of this class, i.e. we duplicate the data but still share the code.

*Intelligent Backtracking.* Pipelining makes it easy to implement optimizations such as existential optimization and intelligent backtracking (Chimenti *et al.*, 1990). Take for instance the following example:

### Example 3.1

#### Intelligent Backtracking

$$\text{query3}(A, B) \leftarrow \text{b1}(A), \text{p}(A, B), \text{b2}(A).$$

Take the situation where the first A-value generated by  $\text{b1}$  is passed to  $\text{p}(A, B)$ , which succeeds and passes the value of A to  $\text{b2}$ . If the first call to this third goal fails, there is no point in going back to  $\text{p}$ , since this can only return a new value for B. Instead, we have to jump back to  $\text{b1}$  for a new value of A. In an eager approach, all the B-values corresponding to each A are computed, even when they cannot satisfy  $\text{b2}$ .

Similar optimizations were also supported in *LDL* (Chimenti et al., 1990), but with various limitations, e.g. existential optimization was not applied to recursive predicates, since these were not pipelined. In *LDL++*, the techniques are applied uniformly, since pipelining is now used in the computation of all predicates, including recursive ones.

### 3.3 External databases

A most useful feature of the *LDL++* system is that it supports convenient and efficient access to external databases. As shown in figure 1, the External Database Interface (EDI) provides the capability to interact with external databases. The system is equipped with a generic SQL interface as well as an object-oriented design that allows easy access to external database systems from different vendors. To link the system with a specific external database, it is only necessary to write a small amount of code to implement vendor-specific drivers to handle data conversion and local SQL dialects. The current *LDL++* system can link directly with Sybase, Oracle, DB2 and indirectly with other databases via JDBC<sup>3</sup>.

The rules in a program make no distinction between internal and external relations. Relations from external SQL databases are declared in the *LDL++* schema just like internal relations, with the additional specification of the type and the name of the SQL server holding the data. As a result, these external resources are transparent to the inference engine, and applications can access different databases without changes. The EDI can also access data stored in files.

The following example shows the *LDL++* schema declarations used to access an external relation `employee` in the database `payroll` running on the server `sybase_tarski`.

#### Example 3.2

Schema Declaration to external Sybase server.

```
database({
    sybase::employee(NAME:char(30),SALARY:int, MANAGER:char(30))
        from sybase_tarski
        use payroll
        user_name 'john'
        application_name 'downsizing'
        interface_filename '/tmp/ldl++/demo/interfaces'
        password nhoj
}).
```

The *LDL++* system generates SQL queries that off-loads to the external database server the computation of (i) the join, select, project queries corresponding to positive rule goals, (ii) the set differences corresponding to the negated goals, and (iii) the aggregate operations specified in the heads of the rules.

<sup>3</sup> Sybase is a trademark of Sybase Inc., Oracle is a trademark of Oracle Inc., DB2 is a trademark of IBM Inc.



In the following example the rule defines expensive employees as those who make over 75,000 and more than their managers:

*Example 3.3*

SQL Generation

```
expensive_employee(Name) <-
    employee(Name, Salary1, Manager),
    Salary1 > 75000,
    employee(Manager, Salary2, _),
    Salary1 > Salary2.
```

The *LDL++* compiler collapses all the goals of this rule and transforms it into the following SQL node:

```
expensive_employee(Name) <- sql_node(Name).
```

where `sql_node` denotes the following SQL query sent to external database server:

```
SELECT  employee_0.NAME
FROM    employee employee_0, employee employee_1
WHERE   employee_0.SALARY > 75000 AND
        employee_1.NAME = employee_0.MANAGER AND
        employee_0.SALARY > employee_1.SALARY
```

Consequently, access to the external database via *LDL++* is as efficient as for queries written directly in SQL. Rules with negated goals are also supported and implemented via the NOT EXIST construct of SQL. The *LDL++* SQL interface also supports updates to external databases, including set-oriented updates with qualification conditions. Updates to external relations follow the same syntax and semantics as the updates to local relations. The execution of each query form is viewed as a new transaction: either it reaches its commit point or the transaction is aborted.

To better support middleware applications, the coupling of *LDL++* with external databases was further enhanced as follows:

- *Literal collapsing*: the goals in the body of a rule are reordered to ensure that several goals using database relations can now be supported as a single SQL subquery to be offloaded to the DBMS.
- *Rule compression*: to offload more complex and powerful queries the remote database, literals from multiple levels of rules are combined and the rules are compressed vertically.
- *Aggregates*: rules that contain standard SQL aggregates in their heads can also be offloaded to the remote SQL system.

### 3.4 Procedural language interface

As shown in figure 1, the *LDL++* system is designed to achieve an open architecture where links with procedural languages, such C/C++, can be established in two ways:

- via the Application Programming Interface (API) which allows applications to drive the system; and
- via the External Predicate Manager which allows C/C++ functions to be imported into the inference engine as external predicates.

Via the API, any C/C++ routine can call the *LDL++* inference engine. The API provides a set of functions that enable applications to instruct the *LDL++* engine to load a schema, load rules, compile query forms, send queries and retrieve results.

Via the external predicate manager, function defined in C/C++ can be imported into *LDL++* and treated as logical predicates callable as rule goals. A library of C/C++ functions is also provided to facilitate the manipulation of internal *LDL++* objects, and the return of multiple answers by the external functions. Therefore, external functions can have the same behavior as internal predicates in all aspects, including flow of control and backtracking. Details on these interfaces can be found in Zaniolo *et al.* (1998).

#### 4 Applications

The deployment of the *LDL* and *LDL++* prototypes in various real-life applications have much contributed to understanding the advantages and limitations of deductive databases in key application domains (Tsur, 1990a, 1990b). Moreover, this experience with application problems, has greatly influenced the design of the *LDL++* system and its successive improvements.

*Recursive Queries.* Our first focus was to compute transitive closures and to solve various graph problems requiring recursive queries, such as Bill-of-Materials (Zaniolo *et al.*, 1997). Unfortunately, many of these applications also require that set-aggregates, such as counts and minima, be computed during the recursive traversal of the graph. Therefore, these applications could not be expressed in *LDL* which only supported stratified semantics, and thus disallowed the use of negation and aggregation within recursive cliques. Going beyond stratification thus became a major design objective for *LDL++*.

*Rapid Prototyping of Information Systems.* Rapid prototyping from E-R specifications has frequently been suggested as the solution for the productivity bottleneck in information system design. Deductive databases provide a rule-based language for encoding executable specifications, that is preferable to Prolog and 4GL systems used in the past, because their completely declarative semantics provides a better basis for specifications and formal methods. Indeed, *LDL* proved to be the tool of choice in the rapid *Prototyping of Information Systems* in conjunction with a structured-design methodology called POS (Process, Object and State) (Ackley *et al.*, 1990; Tryon, 1991). Our proof-of-concept experiment confirmed the great potential of deductive databases for the rapid prototyping of information systems; but this also showed the need for a richer environment that also supports prototyping of

graphical interfaces, and the use of E-R based CASE tools. A large investment in producing such tools is probably needed before this application area can produce a commercial success for deductive databases.

*Middleware.* At MCC, *LDL++* was used in the CARNOT/INFOSLEUTH project to support semantic agents that carry out distributed, coordinated queries over a network of databases (Ong *et al.*, 1995). In particular, *LDL++* was used to implement the ontology-driven mapping between different schemas; the main functions performed by *LDL++* include (i) transforming conceptual requests by users into a collection of cooperating queries, (ii) performing the needed data conversion, and (iii) offloading to SQL statements executable on local schemas (for both relational and O-O databases).

*Scientific Databases.* The *LDL++* system provided a sound environment on which to experiment with next-generation database applications, e.g. to support domain science research, where complex data objects and novel query and inferencing capabilities are required.

A first area of interest was molecular biology, where several pilot applications relating to the Human Genome initiative (Erickson, 1992) were developed (Overbeek *et al.*, 1990; Tsur *et al.*, 1990). *LDL++* rules were also used to model and support taxonomies and concepts from the biological domain, and to bridge the gap between high-level scientific models and low-level experimental data when searching and retrieving domain information (Tsur, 1990b).

A second research area involves geophysical databases for atmospheric and climatic studies (Muntz *et al.*, 1995). For instance, there is a need for detecting and tracking over time and space the evolution of synoptic weather patterns, such as cyclones. The use of *LDL++* afforded the rapid development of queries requiring sophisticated spatio-temporal reasoning on the geographical database. This first prototype was then modified to cope with the large volume of data required, by off-loading much of the search work to the underlying database. Special constructs and operators were also added to express cyclone queries (Muntz *et al.*, 1995).

*Knowledge Discovery and Decision Support Applications.* The potential of the *LDL++* technology in this important application area was clear from the start (Naqvian and Tsur, 1989), when our efforts concentrated on providing the analyst with powerful tools for the verification and refinement of scientific hypotheses (Tsur, 1990a). In our early experiments, the expert would write complex verification rules that were then applied to the data. *LDL++* proved well-suited for the rapid prototyping of these rules, yielding what became known as the 'data dredging' paradigm (Tsur, 1990a).

A more flexible methodology was later developed combining the deductive rules with inductive tools, such as classifiers or Bayesian estimation techniques. A prototype of a system combining both the deductive and inductive methods is the 'Knowledge Miner' (Shen *et al.*, 1994), which was used in the discovery of rules from

a database of chemical process data;  $\mathcal{L}\mathcal{D}\mathcal{L}++$  meta predicates proved very useful in this experiment (Shen et al., 1996).

Other experiments demonstrated the effectiveness of the system in performing important auxiliary tasks, such as data cleaning (Shaw et al., 1993; Sheth et al., 1995). In these applications, the declarative power of  $\mathcal{L}\mathcal{D}\mathcal{L}++$  is used to specify the rules that define correct data. These allow record-by-record verification of data for correctness but also the identification of *sets* of records, whose combination violates the integrity of the data. Finally, the rules are used to clean (i.e. correct) inconsistent data. This capability can either be used prior to the loading of data into the database, or during the updating of the data after loading. This early investigations paved the way for a major research project discussed next focusing on using  $\mathcal{L}\mathcal{D}\mathcal{L}++$  in datamining applications.

*Developing Data Mining Applications.* The results of extensive experiences with an  $\mathcal{L}\mathcal{D}\mathcal{L}++$  based environment for knowledge discovery were reported in Giannotti et al. (1999) and Bonchi et al. (1999). The first study (Giannotti et al., 1999) describes the experience with a fraud detection application, while the second one reports on a marketing application using market basket analysis techniques (Bonchi et al., 1999). In both studies,  $\mathcal{L}\mathcal{D}\mathcal{L}++$  proved effective at supporting the many diverse steps involved in the KDD process. In Bonchi et al. (1999), the authors explain the rationale for their approach and the reasons for their success, by observing that the process of making decisions requires the integration of two kinds of activities: (i) knowledge acquisition from data (inductive reasoning), and (ii) deductive reasoning about the knowledge thus induced, using expert rules that characterize the specific business domain. Activity (i) relies mostly on datamining functions and algorithms that extract implicit knowledge from raw data by performing aggregation and statistical analysis on the database. A database-oriented rule-based system, such as  $\mathcal{L}\mathcal{D}\mathcal{L}++$ , is effective at driving and integrating the different tasks involved in (i) and very effective in activity (ii) where the results of task (i) are refined, interpreted and integrated with domain knowledge and business rules characterizing the specific application.

For instance, association rules derived from market basket analysis are often too low-level to be directly used for marketing decisions. Indeed, market analysts seek answers to higher-level questions, such as “Is the supermarket assortment adequate for the company’s target customer class?” or “Is a promotional campaign effective in establishing a desired purchasing habit in the target class of customers?”.  $\mathcal{L}\mathcal{D}\mathcal{L}++$  deductive rules were used in Bonchi et al. (1999) to drive and control the overall discovery process and to refine the raw association rules produced by datamining algorithms into knowledge of interest to the business. For instance,  $\mathcal{L}\mathcal{D}\mathcal{L}++$  would be used to express queries such as “Which rules survive/decay as one moves up or down the product hierarchy?” or “What rules have been effected by the recent promotions” (Bonchi et al., 1999).

The most useful properties of  $\mathcal{L}\mathcal{D}\mathcal{L}++$  mentioned in these studies (Giannotti et al., 1999; Bonchi et al., 1999; Giannotti et al., 2001a) were flexibility, capability to adapt to the analyst’s needs, and modularity, i.e. the ability to clearly separate the different functional components, and provide simple interfaces for their integration.

In particular, the user defined aggregates described in section 2.2 played a pivotal roles in these datamining applications since datamining functions (performing the inductive tasks) were modelled as user-defined aggregates which could then be conveniently invoked by the  $\mathcal{L}\mathcal{D}\mathcal{L}++$  rules performing the deductive tasks (Giannotti *et al.*, 2001a). The performance and scalability challenge was then addressed by encoding these user-defined aggregates by means of  $\mathcal{L}\mathcal{D}\mathcal{L}++$  procedural extensions, and, for database resident data, offloading critical tasks to the database system containing the data (Giannotti *et al.*, 2001a).

*Lessons Learned.* The original motivations for the development of the original  $\mathcal{L}\mathcal{D}\mathcal{L}$  system was the desire to extend relational query languages to support the development of complete applications, thus eliminating the impedance mismatch from which applications using embedded SQL are now suffering. In particular, data intensive expert systems were the intended ‘killer’ applications for  $\mathcal{L}\mathcal{D}\mathcal{L}$ . It was believed that such applications call for combining databases and logic programming into a rule-based language capable of expressing reasoning, knowledge representation, and database queries. While the original application area failed to generate much commercial demand, other very promising areas emerged since then. Indeed the success of  $\mathcal{L}\mathcal{D}\mathcal{L}++$  in several areas is remarkable, considering that  $\mathcal{L}\mathcal{D}\mathcal{L}++$  is suffering from the combined drawbacks of (i) being a research prototype (rather than a supported product), and yet (ii) being subject to severe licensing limitations. Unless the situation changes and these two handicaps are removed, the only opportunities for commercial deployments will come from influencing other systems; i.e. from system that borrow the  $\mathcal{L}\mathcal{D}\mathcal{L}++$  technology to gain an edge in advanced application areas, such as datamining and decision support systems.

## 5 Conclusion

Among the many remarkable projects and prototypes (Ramakrishnan and Ullman, 1995) developed in the field of logic and databases (Minker, 1996), the  $\mathcal{L}\mathcal{D}\mathcal{L}/\mathcal{L}\mathcal{D}\mathcal{L}++$  project occupies a prominent position because the level and duration of its research endeavor, which brought together theory, systems, and applications. By all objective measures, the  $\mathcal{L}\mathcal{D}\mathcal{L}++$  project succeeded in its research objectives. In particular, the nondeterministic and nonmonotonic constructs now supported in  $\mathcal{L}\mathcal{D}\mathcal{L}++$  take declarative logic-based semantics well beyond stratification in terms of power and expressivity (and stratified negation is already more powerful than SLD-NF). The  $\mathcal{L}\mathcal{D}\mathcal{L}++$  system supports well the language and its applications. In particular, the pipelined execution model dovetails with constructs such as choice and aggregates (and incremental answer generation), while the system’s open architecture supports tight coupling with external databases, JDBC, and other procedural languages. The merits of the  $\mathcal{L}\mathcal{D}\mathcal{L}++$  technology, and therefore of deductive databases in the large, have been demonstrated in several pilot applications – particularly datamining applications.

Although there is no current plan to develop  $\mathcal{L}\mathcal{D}\mathcal{L}++$  commercially, there remain several exciting opportunities to transfer its logic-oriented technology to

related fields. For instance, the new query and data manipulation languages for web documents, particularly XML documents, bear affinity to logic-based rule languages. Another is the extension to SQL databases of the new constructs and non-stratified semantics developed for  $\mathcal{L}\mathcal{D}\mathcal{L}++$ : in fact, the use of monotonic aggregates in SQL has already been explored in Wang and Zaniolo (2000).

### Appendix I: Aggregates in logic

The expressive power of choice can be used to provide a formal definition of aggregates in logic. Say for instance that we want to define the aggregate `avg` that returns the average of all  $Y$ -values that satisfy  $d(Y)$ . By the notation used in  $\mathcal{L}\mathcal{D}\mathcal{L}$  (Chimenti et al., 1990), *CORAL* (Ramakrishnan et al., 1993) and  $\mathcal{L}\mathcal{D}\mathcal{L}++$ , this computation can be specified by the following rule:

$$p(\text{avg}(Y)) \leftarrow d(Y).$$

A logic-based equivalent for this rule is

$$p(Y) \leftarrow \text{results}(\text{avg}, Y).$$

where `results(avg, Y)` is derived from  $d(Y)$  by (i) the chain rules, (ii) the `cagr` rules and (iii) the return rules.

The chain rules are those of Example 2.3 that place the elements of  $d(Y)$  into an order-inducing chain.

$$\begin{aligned} &\text{chain}(\text{nil}, \text{nil}). \\ \text{chain}(X, Y) &\leftarrow \text{chain}(\_, X), d(Y), \\ &\quad \text{choice}((X), (Y)), \text{choice}((Y), (X)). \end{aligned}$$

Now, we can define the `cagr` rules to perform the inductive computation by calling the `single` and `multi` rules as follows:

$$\begin{aligned} \text{cagr}(\text{AgName}, Y, \text{New}) &\leftarrow \text{chain}(\text{nil}, Y), Y \neq \text{nil}, \text{single}(\text{avg}, Y, \text{New}). \\ \text{cagr}(\text{AgName}, Y2, \text{New}) &\leftarrow \text{chain}(Y1, Y2), \text{cagr}(\text{AgName}, Y1, \text{Old}), \\ &\quad \text{multi}(\text{AgName}, Y2, \text{Old}, \text{New}). \end{aligned}$$

Thus, the `cagr` rules are used to memorize the previous results, and to apply (i) `single` to the first element of  $d(Y)$  (i.e. for the pattern `chain(nil, Y)`) and (ii) `multi` to the successive elements. The return rules are as follows:

$$\begin{aligned} \text{results}(\text{AgName}, \text{Yield}) &\leftarrow \text{chain}(Y1, Y2), \text{cagr}(\text{AgName}, Y1, \text{Old}), \\ &\quad \text{ereturn}(\text{AgName}, Y2, \text{Old}, \text{Yield}). \\ \text{results}(\text{AgName}, \text{Yield}) &\leftarrow \text{chain}(X, Y), \neg \text{chain}(Y, \_), \\ &\quad \text{cagr}(\text{AgName}, Y, \text{Old}), \\ &\quad \text{freturn}(\text{AgName}, Y, \text{Old}, \text{Yield}). \end{aligned}$$

Therefore, we first compute `chain`, and then `cagr` that applies the `single` and `multi` rules to every element in the chain. Concurrently, the first `results` rule produces all the results that can be generated by the application of the `ereturn`

rules to the element in the chain. The final returns are instead computed by the second `results` rule that calls on the `freturn` rules once the last element in the chain (i.e. the element without successors) is detected. The second `results` rule is the only rule using negation; in the absence of `freturn` this rule can be removed yielding a positive choice program that is monotonic by Theorem 2. Thus, every aggregate with only early returns is *monotonic with respect to set containment* and can be used freely in recursive rules.

### Acknowledgements

The authors are grateful to the referees for many suggested improvements. This work was partially supported by NSF Grant IIS-0070135.

### References

- Abiteboul, S., Hull, R. and Vianu, V. (1995) *Foundations of Databases*. Addison-Wesley, Reading, MA, 1995.
- Ackley, D., Carasik, R. P., Soon, T., Tyron, D., Tsou, E., Tsur, S. and Zaniolo, C. (1990) System analysis for deductive database environments: an enhanced role for aggregate entities. *Proc. 9th Int. Conference on Entity-Relationship Approach*, pp. 129–142. Lausanne, Switzerland.
- Arni, N., Greco, S. and Saccà, D. (1996) Matching of bounded set terms in the logic language  $\mathcal{L}\mathcal{D}\mathcal{L}++$ . *J. Logic Program* **27**(1), 73–87.
- Bancilhon, F., Maier D., Sagiv, Y. and Ullman, J. (1986) Magic sets and other strange ways to implement logic programs. *Proc. SIGACT-SIGMOD Principles of Database Systems Conference (PODS)*, pp. 1–16.
- Baudinet, M., Chomicki, J. and Wolper, P. (1994) Temporal deductive databases. In: Tansel, A. U., Clifford, J., Gadia, S. K., Jajodia, S. A. and Snodgrass, R. T. (eds.), *Temporal Databases: Theory, Design, and Implementation*, pp. 294–320. Benjamin/Cummings.
- Bonchi, F., Giannotti, F., Mainetto, G. and Pedreschi, D. (1999) Applications of  $\mathcal{L}\mathcal{D}\mathcal{L}++$  to datamining: a classification-based methodology for planning audit strategies in fraud detection. *Proc. 5th ACM SIGKDD Int. Conference on Knowledge Discovery and Data Mining (KDD'99)*, pp. 175–184. ACM.
- Brogi, A., Subrahmanian, V. S. and Zaniolo, C. (1997) The logic of totally and partially ordered plans: a deductive database approach. *Ann. Math. Artif. Intell.* **19**(1–2), 27–58.
- Byrd, L. (1980) Understanding the control flow of Prolog programs. *Proc. Logic Programming Workshop*, pp. 127–138. Debrecen, Hungary.
- Chimenti, D., Gamboa, R. and Krishnamurthy, R. (1989) Abstract machine for  $\mathcal{L}\mathcal{D}\mathcal{L}$ , *Proc. EDBT Conference*, pp. 271–293.
- Chimenti, D., Gamboa, R., Krishnamurthy, R., Naqvi, S. A., Tsur, S. and Zaniolo, C. (1990) The  $\mathcal{L}\mathcal{D}\mathcal{L}$  system prototype. *IEEE J. Data & Knowl. Eng.* **2**(1), 76–90.
- Erickson, D. (1992) Hacking the genome. *Sci. Am.* **266**(4), 128–137.
- Finkelstein, S. J., Mattos, N., Mumick, I. and Pirahesh, J. (1996) Expressing Recursive Queries in SQL. ISO WG3 report X3H2-96-075, March.
- Gelfond, M. and Lifschitz, V. (1988) The stable model semantics of logic programming. *Proc. 5th Int. Conference on Logic Programming*, pp. 1070–1080.
- Giannotti, F., Pedreschi, D., Saccà, D. and Zaniolo, C. (1991) Nondeterminism in deductive databases. *Proc. 2nd Int. Conference on Deductive and Object-Oriented Databases*, pp. 129–141.

- Giannotti, F., Manco, G., Nanni, M. and Pedreschi, D. (1998) On the effective semantics of nondeterministic, nonmonotonic, temporal logic databases. *Computer Science Logic, 12th International Workshop (CSL 1998): Lecture Notes in Computer Science 1584*, pp. 58–72. Springer-Verlag.
- Giannotti, F., Manco, G., Pedreschi, D. and Turini, F. (1999) Experiences with a logic-based knowledge discovery support environment, *ACM SIGMOD Workshop on Research Issues in Data Mining and Knowledge Discovery (DMKD'99)*, Philadelphia, USA.
- Giannotti, F., Manco, G. and Turini, F. (2001) Specifying mining algorithms with iterative user-defined aggregates: a case study. *Principles of Data Mining and Knowledge Discovery, 5th European Conference (PKDD 2001)*, pp. 128–139.
- Giannotti, F., Pedreschi, D. and Zaniolo, C. (2001) Semantics and expressive power of non-deterministic constructs in deductive databases. *J. Comput. & Syst. Sci.* **62**(1), 15–42.
- Greco, S. and Sacca, D. (1997) NP Optimization Problems in Datalog. *Logic Programming, Proc. 1997 International Symposium (ILPS 1997)*, pp. 181–195. Port Jefferson, Long Island, NY. MIT Press.
- Han, J. and Kamber, M. (2001) *Data Mining, Concepts and Techniques*, Morgan Kaufman.
- Hellerstein, J. M., Haas, P. J. and Wang, H. J. (1997) Online aggregation. *Proc. ACM-SIGMOD Conference on Management of Data*, pp. 171–182.
- Kemp, D., Ramamohanarao, K. and Stuckey, P. (1995) ELS programs and the efficient evaluation of non-stratified programs by transformation to ELS. In: Ling, T. W., Mendelzon, A. O. and Vieille, L. (eds.), *Proc. 4th Int. Conference on Deductive and Object-Oriented Databases (DOOD'95)*, pp. 91–108. Springer-Verlag.
- Kemp, D. and Ramamohanarao, K. (1998) Efficient recursive aggregation and negation in deductive databases. *Trans. Knowl. & Data Eng.* **10**(5), 727–745.
- Krishnamurthy, R. and Naqvi, S. (1988) Non-deterministic choice in Datalog. *Proc. 3rd International Conference on Data and Knowledge Bases*.
- Lausen, G., Ludäscher, B. and May, W. (1998) On active deductive databases: the Statelog approach. In: Freitag, B., Decker, H., Kifer, M. and Voronkov, A. (eds.), *Transactions and Change in Logic Databases: Lecture Notes in Computer Science 1472*, pp. 69–106. Springer-Verlag.
- Lausen, G., Ludäscher, B. and May, W. (1998) On logical foundations of active databases. In: Chomicki, J. and Saake, G. (eds.), *Logics for Databases and Information Systems*, pp. 375–398. Kluwer Academic.
- Marek, W. and Truszczyński, M. (1991) Autoepistemic logic. *J. ACM*, **38**(3), 588–619.
- Minker, J. (1996) Logic and databases: a 20 year retrospective. In: Pedreschi, D. and Zaniolo, C. (eds.), *Proc. International Workshop on Logic in Databases (LID'96)*, pp. 5–52. Springer-Verlag.
- Muntz, R. R., Shek, E. C. and Zaniolo, C. (1995) Using  $\mathcal{L}\mathcal{D}\mathcal{L}++$  for spatio-temporal reasoning in atmospheric science. In: Ramakrishnan, R. (ed.), *Applications of Logic Databases*, pp. 101–118. Kluwer.
- Naqvi, S. and Tsur, S. (1989) *A Logical Language for Data and Knowledge Bases*. W. H. Freeman.
- Ong, K., Arni, N., Tomlinson, C., Unnikrishnan, C. and Woelk, D. (1995) A deductive database solution to intelligent information retrieval from legacy databases. *Proc. 4th Int. Conference on Database Systems for Advanced Applications (DASFAA 1995)*, pp. 172–179.
- Overbeek, R., Price, M. and Tsur, S. (1990) Automated Interpretation of Genetic Sequencing Gels. MCC Technical Report.
- Phipps, G., Derr, M. and Ross, K. (1991) Glue-Nail: a deductive database system. *Proc. 1991 ACM-SIGMOD Conference on Management of Data*, pp. 308–317.



- Przymusiński, T. (1988) On the declarative and procedural semantics of stratified deductive databases. In: Minker, J. (ed.), *Foundations of Deductive Databases and Logic Programming*, pp. 193–216. Morgan Kaufman.
- Ramakrishnan, R., Srivastava, D. and Sudarshan, S. (1992) CORAL – Control, Relations and Logic. *Proc. 18th VLDB Conference*, pp. 238–250. Morgan Kaufmann.
- Ramakrishnan, R., Srivastava, D., Sudarshan, S. and Seshadri, P. (1993) Implementation of the CORAL deductive database system. *Proc. International ACM SIGMOD Conference on Management of Data*, pp. 167–176.
- Ramakrishnan, R. and Ullman, J. D. (1995) A survey of deductive database systems. *J. Logic Program.* **23**(2), 125–149.
- Ross, K. A. (1994) Modular stratification and magic sets for Datalog programs with negation. *J. ACM*, **41**(6), 1216–1266.
- Ross, K. A. and Sagiv, Y. (1997) Monotonic aggregation in deductive database. *J. Comput. & Syst. Sci.* **54**(1), 79–97.
- Saccà, D. and Zaniolo, C. (1990) Stable models and nondeterminism in logic programs with negation. *Proc. 9th ACM SIGACT-SIGMOD-SIGART Symposium on Principles of Database Systems*, pp. 205–218.
- Schlipf, J. S. (1993) A survey of complexity and undecidability results in logic programming. *Proc. Workshop on Structural Complexity and Recursion-Theoretic Methods in Logic Programming*, pp. 143–164.
- Shaw, S., Foggiato-Bish, L., Garcia, I., Tillman, G., Tryon, D., Wood, W. and Zaniolo C. (1993) Improving data quality via  $\mathcal{L}\mathcal{D}\mathcal{L}++$ . *ILP'93 Workshop on Programming with Logic Databases*, Vancouver, Canada.
- Shen, W., Mitbander, W., Ong, K. and Zaniolo, C. (1994) Using metaqueries to integrate inductive learning and deductive database technology. *KDD Workshop 1994: Knowledge Discovery from Databases*, pp. 335–346. Seattle, WA. AAAI Press.
- Shen, W., Ong, K., Mitbander, B. G. and Zaniolo, C. (1996) Metaqueries for data mining. In: Fayyad, U. M., Piatetsky-Shapiro, G., Smyth, P. and Uthurusamy, R. (eds.), *Advances in Knowledge Discovery and Data Mining*, pp. 201–217. MIT Press.
- Sheth, A. P., Wood, C. and Kashyap, V. (1995) Q-Data: using deductive database technology to improve data quality. In: Ramakrishnan, R. (ed.), *Applications of Logic Databases*, pp. 23–56. Kluwer.
- Shmueli, O., Tsur, S. and Zaniolo, C. (1988) Rewriting of rules containing set terms in a logic database language (LDL). *Proc. 7th ACM Symposium on Principles of Database Systems*, pp. 15–28.
- Van Gelder, A. (1993) Foundations of aggregations in deductive databases. In: Ceri, S., Tanaka, K. and Tsur, S. (eds.), *Proc. Int. Conf. On Deductive and Object-Oriented Databases (DOOD'93)*, pp. 13–34. Springer-Verlag.
- Tryon, D. (1991) Deductive computing: living in the future. *Proc. Monterey Software Conference*.
- Tsur, S., Olken, F. and Naor, D. (1990) Deductive databases for genomic mapping. *Proc. NACL90 Workshop on Deductive Databases*, Austin, TX.
- Tsur, S. (1990) Deductive databases in action. *Proc. 10th ACM SIGACT-SIGMOD-SIGART Symposium on Principles of Database Systems*, pp. 205–218.
- Tsur S. (1990) Data dredging. *Data Eng.* **13**(4), 58–63.
- Ullman, J. D. (1989) *Database and Knowledge-Based Systems, Vols. I and II*. Computer Science Press, Rockville, MD.
- Wang, H. and Zaniolo, C. (2000) User defined aggregates in object-relational systems. *Proc. 16th International Conference on Data Engineering, ICDE'2000*, pp. 135–144. IEEE.

- Zaniolo, C., Arni, N. and Ong, K. (1993) Negation and aggregates in recursive rules: the  $\mathcal{L}\mathcal{D}\mathcal{L}++$  approach. *Proc. 3rd Int. Conference on Deductive and Object-Oriented Databases (DOOD 1993)*, pp. 204–221. Springer-Verlag.
- Zaniolo, C. (1994) A unified semantics for active and deductive databases. In: Paton, N. W. and Williams, H. (eds.), *Proc. Workshop on Rules in Database Systems (RIDS93)*, pp. 271–287. Springer-Verlag.
- Zaniolo, C., Ceri, S., Faloutsos, C., Snodgrass, R. T., Subrahmanian, V. S. and Zicari, R. (1997) *Advanced Database Systems*, Morgan Kaufmann.
- Zaniolo, C. (1997) The nonmonotonic semantics of active rules in deductive databases. *Proc. 5th Int. Conference on Deductive and Object-Oriented Databases (DOOD 1997)*, pp. 265–282. Springer-Verlag.
- Zaniolo, C., Tsur, S. and Wang, H. (1998)  $\mathcal{L}\mathcal{D}\mathcal{L}++$  Documentation and Web Demo. <http://www.cs.ucla.edu/ldl>.
- Zaniolo, C. and Wang, H. (1999) Logic-based user-defined aggregates for the next generation of database systems. In: Apt, K. R., Marek, V., Truszczyński, M. and Warren, D. S. (eds.), *The Logic Programming Paradigm: Current Trends and Future Directions*, pp. 401–424. Springer-Verlag.