Implementation of the CORAL Deductive Database System*

Raghu Ramakrishnan†
University of Wisconsin, Madison

Divesh Srivastava
University of Wisconsin, Madison

S. Sudarshan‡
AT&T Bell Labs, Murray Hill.

Praveen Seshadri
University of Wisconsin, Madison

Abstract

CORAL is a deductive database system that supports a rich declarative language, provides a wide range of evaluation methods, and allows a combination of declarative and imperative programming. The data can be persistent on disk or can reside in main-memory. We describe the architecture and implementation of CORAL.

There were two important goals in the design of the CORAL architecture: (1) to integrate the different evaluation strategies in a reasonable fashion, and (2) to allow users to influence the optimization techniques used so as to exploit the full power of the CORAL implementation. A CORAL declarative program can be organized as a collection of interacting modules and this modular structure is the key to satisfying both these goals. The high level module interface allows modules with different evaluation techniques to interact in a transparent fashion. Further, users can optionally tailor the execution of a program by selecting from among a wide range of control choices at the level of each module.

CORAL also has an interface with C++, and users can program in a combination of declarative CORAL, and C++ extended with CORAL primitives. A high degree of *extensibility* is provided by allowing C++ programmers to use the class structure of C++ to enhance the CORAL implementation.

*This research was supported by a David and Lucile Packard Foundation Fellowship in Science and engineering, a Presidential Young Investigator Award, with matching grants from Digital Equipment Corporation, Tandem and Xerox, and NSF grant IRI-9011563.

†The addresses of the authors are Computer Sciences Department, University of Wisconsin, Madison, WI 53706, USA, and AT&T Bell Laboratories, 600 Mountain Avenue, Murray Hill, NJ 07974, USA. The authors email addresses are {raghu,divesh,praveen}@cs.wisc.edu and sudarsha@research.att.com.

‡The work of this author was performed largely while he was at the University of Wisconsin, Madison, and was partially supported by the grants listed above.

0

1 Introduction

In this paper, we discuss the design and implementation of the CORAL deductive database system. CORAL combines features of a database query language, such as efficient treatment of large relations, aggregate operations and declarative semantics, with those of a logic programming language, such as more powerful inference capabilities and support for structured data. Support for persistent relations is provided using the EXODUS storage manager [2]. A unique feature of CORAL is that it provides a wide range of query evaluation strategies (top-down evaluation and several variants of bottom-up evaluation) and allows users to optionally tailor execution of a program through high-level annotations. Applications in which large amounts of data must be extensively analyzed are likely to benefit from this combination of features. In comparison to other deductive database systems such as Aditi [28], EKS-V1 [29], LDL [27], LOLA [4] and Nail-Glue [11], CORAL provides a more powerful language and supports a much wider range of optimization techniques.

We highlight several design decisions that allowed us to integrate diverse evaluation techniques and optimizations in a nearly seamless fashion. Specifically, we consider the following issues:

1. Data representation (e.g. constants, lists, sets).
2. Relation representation and implementation (e.g. main-memory and disk-resident).
3. Index structures (e.g. hash-structures and B-trees).
4. Evaluation techniques (e.g. materialization and pipelining)

In the CORAL implementation, we divide evaluation into a number of distinct subtasks, and provide a clean interface between the subtasks; relevant optimization techniques can be (almost) independently applied to each subtask. Extensibility in database systems has received much attention lately, and we believe that the CORAL experience offers guidelines for addressing several important issues.

One of our goals was to allow users to exploit the full

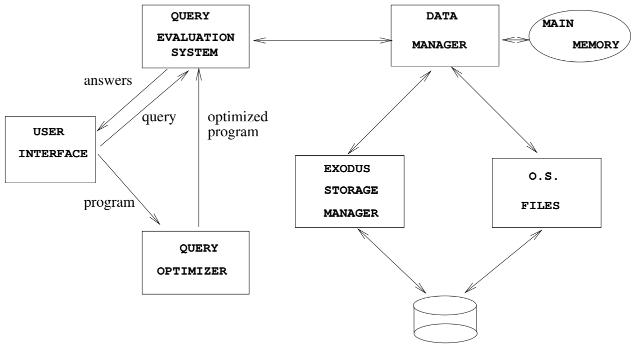

Figure 1: CORAL System Architecture

power of the implementation. CORAL supports a very rich language, and we believe that some user guidance is critical to effectively optimizing many sophisticated programs. The problem is to provide users with the ability to choose from the suite of optimizations supported by CORAL in a relatively orthogonal and high-level way, and to use a combination of optimizations for different parts of a program. The *module* structure was the key to solving this problem. The interface between modules is kept at a high level; evaluation techniques can be chosen on a per-module basis through (optional) *annotations*, and modules with different evaluation techniques can interact in a nearly transparent fashion.

An overview of the CORAL declarative language is presented in [22]. The query language supports general Horn clauses with complex terms, set-grouping, aggregate operations, negation and data that contains universally quantified variables. The details of the language are outside the scope of this paper. Many features of the implementation ranging from low-level structures to the interactive system environment have also been omitted due to shortage of space, and are described in the full version of the paper. In Sections 2–5, we discuss some important aspects of the system implementation. Section 2 contains an overview of the CORAL system architecture, Section 3 describes the underlying representation of the data and Section 4 provides an overview of query evaluation and optimization. Section 5 is the main section that deals with implementation issues. It covers the basic strategies used in evaluating a module, as well as several important refinements. It also addresses user guidance of query optimization, and the interaction in the evaluation of different modules. The CORAL/C++ interface and support for extensibility in CORAL, including the addition of new data types and operations, and new relation and index implementations, are discussed in Sections 6 and 7. We discuss related systems in Section 8. Finally, we provide a retrospective discussion of the CORAL design and outline future research directions in Section 9.

2 CORAL System Architecture

The architecture of the CORAL deductive system is shown in Figure 1. Persistent data is stored either in text files, or using the EXODUS storage manager [2], which has a client-server architecture. Each CORAL single-user process is a client that can access the common persistent data from the EXODUS server. Multiple CORAL processes could interact by accessing persistent data stored using the EXODUS storage manager. Transactions and concurrency control are supported by the EXODUS toolkit, and thus by CORAL. However, within each CORAL process, all data that is *not* managed by the EXODUS storage manager is strictly local to the process. Most of the effort of design and implementation in CORAL has concentrated on the single-user client, and the implementation has focused on operation out of main memory. CORAL incorporates several optimizations, such as using pointers to avoid copying data, that are useful when operating out of main-memory. However, it is important to note that the basic design of the CORAL system does not assume operation out of main-memory, and can use join algorithms and access methods tailored to disk-resident data.

Data stored in text files can be 'consulted', at which point the data is converted into main-memory relations. Data stored using the EXODUS storage manager is paged into client EXODUS buffers on demand, making use of the indexing and scan facilities of the storage manager.

The query processing system consists of two main parts — a query optimizer and a query evaluation system. Simple queries (selecting facts from a base relation, for instance) can be typed in at the user interface. Complex queries are typically defined in declarative 'program modules' that export predicates (views) with associated 'query forms' (i.e. specifications of what kinds of queries, or selections, are allowed on the predicate).

The query optimizer takes a program module and a query form as input, and generates a rewritten program that is optimized for the specified query forms. In addition to doing rewriting transformations, the optimizer adds several *control annotations* (to those, if any, specified by the user). The rewritten program is converted into an internal representation that is used by the query evaluation system.

The query evaluation system takes as input annotated declarative programs (in an internal representation), and database relations. The annotations in the declarative programs provide execution hints and directives. The query evaluation system *interprets* the internal form of the optimized program. We also developed a *compiled* version of CORAL, that generated a C++ program from each user program. (This is the approach taken by LDL [13, 3].) We found that this approach took a much longer time to compile programs (compilation in LDL is also quite slow). Moreover, the resulting gain in execution speed in CORAL was minimal since declarative CORAL currently is not strongly typed, and data manipulation is essentially interpretive. (The issue is discussed further in Section 9.) We have therefore focused on the interpreted version; 'consulting' a program takes very little time and is comparable to Prolog systems, which makes CORAL convenient for interactive program development.

The query evaluation system has a well defined 'get-next-tuple' interface with the data manager for access to relations. This interface is independent of how the relation is defined (as a base relation, declaratively through rules, or through system- or user-defined C++ code). In conjunction with the modular nature of the CORAL language, such a high level interface is very useful, since it allows the different modules to be evaluated using different strategies. It is important to stress that CORAL *does* manipulate data in a set-oriented fashion, and the 'get-next-tuple' interface is merely an abstraction provided to support modularity in the language.

CORAL supports an interface to C++, extended with several features that provide the abstraction of relations and tuples. C++ can be used to define new relations as well as manipulate relations computed using embedded declarative CORAL rules. The CORAL/C++ interface is important for the development of large applications.

3 The Data Manager

The data manager (DM) is responsible for maintaining and manipulating the data in relations. In discussing the DM, we also discuss the representation of the various data types. While the representation of simple types is straight-forward, complex structural types and variables (used to model incomplete data) present interesting challenges. The efficiency with which such data can be processed depends in large part on the manner in which it is represented in the system. This section therefore presents the data representation at a fairly detailed level, and this facilitates the discussion of evaluation techniques in the subsequent sections.

The CORAL system is implemented in C++, and all data types are defined as C++ classes. *Extensibility* is an important goal of the CORAL system. In particular, we view support for user-defined abstract data types as important. In order to provide this support, CORAL provides the class Arg that is the root of all CORAL data-types.

3.1 Representation of Terms

The primitive data types provided in the CORAL system include integers, doubles, strings, and arbitrary precision integers.[1] The current implementation restricts data that is stored using the EXODUS storage manager to be limited to terms of these primitive types. Such data is stored on disk in its machine representation, while in memory, the data types are implemented as subclasses of Arg.

The evaluation of rules in CORAL is based on the operation of *unification* that generates bindings for variables based on patterns in the rules and the data. An important feature of the CORAL implementation of data types is the support for unique identifiers to make unification of complex terms very efficient. Such support is critical for efficient declarative program evaluation in the presence of complex terms. In CORAL, each type can define how it generates unique identifiers, independent of how other types construct their unique identifiers; due to this orthogonality, no further integration is needed to generate unique identifiers for terms built using several different kinds of types. This is very important for supporting extensibility and the creation of new abstract data types.

Terms built from a function symbol, or *functor*, are important for representing structured information. The following fact contains a functor term built from a symbol $addr$:
$person$('John', $addr$('1 Main St.', 'NY', 'NY', '10000')).
A functor term models a C++ struct (or a Pascal record); functor terms can be used to construct complex structures such as lists. The $addr(\ldots)$ term above is represented by a record containing (1) the function symbol $addr$, (2) an array of arguments, and (3) extra information to make unification of such terms efficient. The last item is important for large terms such as lists. The current implementation of CORAL uses a modified version of hash-consing ([5]) that operates in a lazy fashion. Hash-consing assigns unique identifiers to each ground (i.e. variable free) functor term, such that two (ground) functor terms unify if and only if their unique identifiers are the same. We note that such identifiers cannot be assigned to functor terms that contain free variables, and these have to be handled differently.

Variables are primitive values in CORAL, since CORAL allows facts (and not just rules) to contain variables; in this, CORAL differs from most other deductive database systems. The semantics of such non-ground facts is that all variables are universally quantified in the fact. Although the basic representation of variables is fairly simple, the representation is complicated by requirements of efficiency when using non-ground facts in rules (see [26]). Suppose we want to make an inference using a rule. Variables in the rule may get bound in the course of an inference. A naive scheme would replace every reference to the variable by its binding. It is more efficient however to record variable bindings in a *binding environment*, at least during the course of an inference. A binding environment (often referred to as a *bindenv*) is a structure that stores bindings for variables. Therefore, whenever a variable is accessed during an inference, a corresponding binding environment must be accessed to determine if the variable has been bound. In Figure 2, we show the representation of the term $f(X, 10, Y)$, where X is bound to 25 and Y is bound to Z, and Z is bound to 50 in a separate bindenv.

3.2 Representation of Relations

CORAL currently supports in-memory hash-relations, as well as persistent relations (the latter using the EXODUS storage manager [2]). Multiple indices can be created on relations, and can be added to existing relations. The relation

[1] Arbitrary precision integers are supported using the BigNum package provided by DEC France.

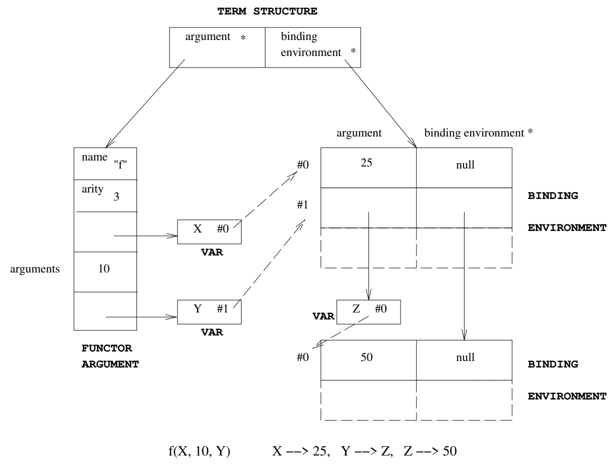

Figure 2: Representation of an Example Term

interface is designed to make the addition of new relation implementations relatively easy.

CORAL relations support the ability to get *marks* into a relation, and distinguish between facts inserted after a mark was obtained and facts inserted before the mark was obtained. This feature is important for the implementation of all variants of semi-naive evaluation [1, 20]. The implementation of this extension involves creating subsidiary relations, one corresponding to each interval between marks, and transparently providing the union of the subsidiary relations corresponding to the desired range of marks. A benefit of this organization is that it can be combined with the indexing mechanisms used for the relation (the indexing mechanisms are used on each subsidiary relation).

CORAL uses the EXODUS storage manager to support persistent relations. Currently, tuples in a persistent relation are restricted to have fields of primitive types only. EXODUS uses a client-server architecture; CORAL is the client process, and maintains buffers for persistent relations. If a requested tuple is not in the client buffer pool, a request is forwarded to the EXODUS server and the page with the requested tuple is retrieved.[2]

3.3 Representation of Indices

Hash-based indices for in-memory relations and B-tree indices for persistent relations are currently available in the CORAL system. CORAL allows the specification of two types of hash-based indices: (1) *argument form indices*, and (2) *pattern form indices*. The first form is the traditional multi-attribute hash index on a subset of the arguments of a relation. All facts that contain a variable in an indexed argument are hashed to a special value. The second form is more sophisticated, and allows us to retrieve precisely those facts that match a specified pattern, where the pattern can contain variables. Such indices are of great use when dealing with complex objects created using functors. One can retrieve, for example, those tuples in relation *append* that have as the first argument a list that matches $[X|[1,2,3]]$. A tuple $([5|[1,2,3]],[4],[5,1,2,3,4])$ would match this pattern (see [24]).

4 Overview of Query Evaluation

A number of query evaluation strategies have been developed for deductive databases, and each technique is particularly efficient for some classes of programs, but may perform relatively poorly on others. It is our premise that in such a powerful language, completely automatic optimization can only be an ideal; the programmer must be able to provide hints or *annotations* and occasionally even override the system's decisions in order to obtain good performance across a wide range of programs. Since annotations can be expressed at a high level, they give the programmer the power to control optimization and evaluation in a relatively abstract manner. A detailed description of the annotations provided by CORAL is found in [22]; we mention some of them when discussing the query evaluation techniques.

The CORAL programmer decides (on a per-module basis) whether to use one of two basic evaluation approaches, namely *pipelining* or *materialization*, which are discussed in Section 5. Many other optimizations depend upon the choice of the basic evaluation mode. The optimizer generates annotations that govern many run-time actions, and, if materialization is chosen, does source-to-source rewriting of the user's program. We discuss these two major tasks of the optimizer below.

4.1 Rewriting Techniques

Materialized evaluation in CORAL is essentially a fixpoint evaluation using bottom-up iteration on the program rules (see, for instance, [23]). If this is done on the original program, selections in a query are not utilized. Several program transformations have been proposed to 'propagate' such selections, and many of these are implemented in CORAL. The desired selection pattern is specified using a query form, where a 'bound' argument indicates that any binding in that argument position of the query is to be propagated.

The default rewriting technique is Supplementary Magic Templates [16]. The rewriting can be tailored to propagate bindings across subgoals in a rule body using different subgoal orderings; CORAL uses a left-to-right ordering within the body of a rule by default. Other selection-propagating rewriting techniques supported in CORAL include Magic

[2] The basic architecture of the CORAL system allows data to be accessed directly from the pages in the EXODUS buffer pool as is traditionally done in database systems. In CORAL, copying of data has largely been replaced by sharing using pointers. Pointers should not refer to data in the EXODUS buffer pool since buffer pool space may get reused. Hence CORAL copies data from the buffer pool into the program heap space. A lot of this copying is unnecessary (namely copying data that is not referred to in any facts that are created). We plan to change the implementation to store primitive data types directly in tuples and thereby avoid the need for copying data except when creating new facts.

Templates [16], Supplementary Magic With GoalId Indexing [24], and Context Factoring [14, 7]. By default, CORAL also applies Existential Query Rewriting [17], which seeks to propagate projections.

4.2 Decisions On Run-time Alternatives

In addition to choosing rewriting techniques for materialized evaluation, the optimizer makes a number of decisions that affect execution. The optimizer analyzes the (rewritten) program, and identifies some evaluation and optimization choices that appear appropriate.

The default fixpoint evaluation strategy is called Basic Semi-Naive evaluation (BSN), but a variant, called Predicate Semi-Naive evaluation (PSN), which is better for programs with many mutually recursive predicates, is also available. With respect to semi-naive evaluation, the optimizer is responsible for: (1) join order selection, (2) index selection, (3) deciding whether to refine the basic nested-loops join with *intelligent backtracking* (see, for instance, [13, 3]). These aspects are discussed in detail in the full version of the paper.

The optimizer also decides on the subsumption checks to be carried out on each relation. The default is to do subsumption checks on all relations. A user can ask that a relation be treated as a *multiset*, with as many copies of a tuple as there are derivations for it in the original program.[3] This semantics is supported by carrying out duplicate checks only on the 'magic' predicates; some version of Magic Templates must be used (see [12]).

5 Module Evaluation Strategies

The evaluation of a declarative CORAL program is divided into a number of distinct sub-computations by expressing the program as a collection of modules. Each module is a unit of compilation and its evaluation strategies are independent of the rest of the program. Modules *export* the predicates that they define; a predicate exported from one module is visible to all other modules, and can be used by them in rules. Since different modules may have widely varying evaluation strategies, some relatively high level interface is required for interaction between modules. The basic approach used by CORAL is outlined here.

During the evaluation of a rule r in module M, if a query is generated on a predicate exported by module N, a call is set up on module N. The answers to this query are used iteratively in rule r; each time a new answer to the query is required, rule r requests a new tuple from the interface to module N. The interface makes no assumptions about the evaluation of the module. Module N may have rules that are evaluated in any of several different ways. The module may choose to cache answers between calls, or choose to recompute answers. All this is transparent to the calling module. Similarly, the evaluation of the called module N makes no assumptions about the evaluation of calling module M. This orthogonality permits the free mixing of different evaluation techniques in different modules in CORAL and is central to how different executions in different modules are combined cleanly.

[3] On non-recursive queries, this semantics is consistent with SQL when duplicate checks are omitted.

Two basic evaluation approaches are supported, namely *pipelining* and *materialization*. We describe them in following sections. Pipelining uses facts 'on-the-fly' and does not store them, at the potential cost of recomputation. Materialization stores facts and looks them up to avoid recomputation.

5.1 Pipelining

For *pipelining*, which is essentially top-down evaluation as in Prolog, the rule evaluation code is designed to work in a co-routining fashion — when rule evaluation is invoked, using the *get-next-tuple* interface, it generates an answer (if there is one) and transfers control back to the consumer of answers (the caller). Control is transferred back to the (suspended) rule evaluation when more answers are desired.

At module invocation, the first rule in the list associated with the queried predicate is evaluated. This could involve recursive calls on other rules within the module (which are also evaluated in a similar pipelined fashion). If the rule evaluation of the queried predicate succeeds, the state of the computation is frozen, and the generated answer is returned. A subsequent request for the next answer tuple results in the reactivation of the frozen computation, and processing continues until the next answer is returned. At any stage, if a rule fails to produce an answer, the next rule in the rule list for the head predicate is tried. When there are no more rules to try, the query on the predicate fails. When the topmost query fails, no further answers can be generated, and the pipelined module execution is terminated.

Pipelining guarantees a particular evaluation strategy, and order of execution. While the program is no longer truly 'declarative', programmers can exploit this guarantee and use imperatively defined predicates that have side-effects.

5.2 Materialization

Several variants of materialized evaluation are supported in CORAL: Basic Semi-Naive, Predicate Semi-Naive [20], Ordered Search [21], and the non-ground fact optimizations described in [26]. The variants of *materialization* are all bottom-up fixpoint evaluation methods. Bottom-up evaluation iterates over a set of rules, repeatedly evaluating them until a fixpoint is reached. In order to perform incremental evaluation of rules across multiple iterations, CORAL uses the semi-naive evaluation technique [1, 20]. This technique consists of a rule rewriting part performed at compile time, which creates versions of rules with *delta relations*, and an evaluation part. (The delta relations contain changes in relations since the last iteration.) The evaluation part evaluates each rewritten rule once in each iteration, and performs some updates to the delta relations at the end of the iteration. An evaluation terminates when an iteration produces no new facts.

The basic join mechanism used in the current implementation of CORAL is nested-loops with indexing, where the indices are automatically generated by the optimizer. In a manner similar to Prolog, CORAL maintains a trail of variable bindings when a rule is evaluated; this is used to undo variable bindings when the nested-loops join considers the next tuple in any loop.

5.3 Module and Rule Data Structures

The compilation of a materialized module generates an internal *module structure* that consists of a list of structures corresponding to the strongly connected components (SCCs) of the module[4], and each SCC structure contains structures corresponding to semi-naive rewritten versions of rules. These *semi-naive rule structures* have fields that specify the arguments of each body literal, and the predicates that they correspond to. Each semi-naive rule also contains evaluation order information, pre-computed backtrack points, and pre-computed offsets into a table of relations. In the case of modules to be evaluated using pipelining, the original rules of the module are stored as above, rather than their semi-naive rewritten versions.

5.4 Module Level Control Choices

At the level of the module, a number of choices exist with respect to the evaluation strategy for the module, and the specific optimizations to be used. We describe the implementation of some of these strategies.

5.4.1 Ordered Search

Ordered Search is an evaluation mechanism that orders the evaluation of generated subgoals in a program and thereby provides an important strategy for handling programs with negation, set-grouping and aggregation, that are left-to-right modularly stratified. Full details of Ordered Search are not presented here, but the reader is referred to [21]. The principle of Ordered Search is that the computation is ordered by 'hiding' subgoals. This is achieved by maintaining a 'context' that stores subgoals in an ordered fashion, and that decides at each stage in the evaluation, which subgoal to make available for use next. The order in which generated subgoals are made available for use is somewhat similar to a top-down evaluation.

From an implementation perspective three main changes have to be made. First, the 'context' has to be maintained, and subgoals inserted into it and deleted from it at appropriate points of time. Second, the rewriting phase must use a version of Magic tailored to Ordered Search — the basic rewriting is modified to introduce 'done' literals guarding negated literals and rules that have grouping and aggregation. Third, the evaluation must add a goal ('magic' fact) to the corresponding 'done' predicate when (and only when) all answers to it have been generated. (The context mechanism is used to determine the point at which a goal is considered done.) These changes ensure that rules involving negation, for example, are not applied until enough facts have been computed so that when we make an inference using such a rule, any fact not present may be assumed false.

[4] An SCC is a maximal set of mutually recursive predicates.

5.4.2 The Save Module Facility

In most cases, facts (other than answers to the query) computed during the evaluation of a module are best discarded at the end of the call to the module to save space (since bottom-up evaluation stores many facts, space is generally at a premium). Module calls provide a convenient unit for discarding intermediate answers. By default, CORAL does precisely this. However, there are some cases where this leads to a significant amount of recomputation. This is especially so in cases where the same subgoal in a module is generated in many different invocations of the module. In such cases, the user can tell the CORAL system to maintain the state of the module (i.e. retain generated facts) across calls to the module, and thereby avoid recomputation; we call this facility the *save module* facility.

In the interest of efficient implementation, we have the following restriction on the use of the save module feature: *if a module uses the save module feature, it should not be invoked recursively.* (Note that the predicates defined in the module can be recursive; this does not cause recursive invocations of the module). From an implementation point of view, the challenge is to ensure that no derivations are repeated *across multiple calls to the module.* This requires significant changes to semi-naive evaluation; while the details are omitted here for lack of space, they can be found in the full version of the paper.

5.4.3 Lazy Evaluation

In the traditional approach to bottom-up evaluation, all answers to a query are computed by iterating over rules until a fixpoint is reached, and then returning all the answers. Lazy evaluation tries to return the answers at the end of every iteration, instead of at the end of computation. Lazy evaluation is implemented by freezing the state of the computation at the end of an iteration, and returning the answer tuples generated in that iteration. The state is stored with the iterator that is created for the query (recall the 'get-next-tuple' iterative interface). The iterator then iterates over the tuples returned, and when it has stepped through all the tuples, it reactivates the 'frozen' computation that it has stored. This reactivation results in the execution of one more iteration of the rules, and the whole process is repeated until an iteration over the rules produces no new tuples.

```
module s_p.
export s_p(bfff, ffff).
@aggregate_selection p(X, Y, P, C) (X, Y) min(C).

s_p(X, Y, P, C) :- s_p_length(X, Y, C), p(X, Y, P, C).
s_p_length(X, Y, min(< C >)) :- p(X, Y, P, C).
p(X, Y, P1, C1) :- p(X, Z, P, C), edge(Z, Y, EC),
        append([edge(Z, Y)], P, P1), C1 = C + EC.
p(X, Y, [edge(X, Y)], C) :- edge(X, Y, C).

end_module.
```

Figure 3: Program Shortest_Path

5.5 Predicate Level Control

CORAL also provides annotations at the level of individual predicates in a module. Annotations to control what indices are created for a predicate were described in Section 3.3. We discuss below a class of predicate-level annotations which we call aggregate selections.

Consider the program Shortest_Path in Figure 3. To compute shortest paths between points, it suffices to use only the shortest paths between pairs of points — path facts that do not correspond to shortest paths are irrelevant. CORAL permits the user to specify an *aggregate selection* on the path predicate p in the manner shown. When a path fact is generated, the aggregate selection causes the system to check if there is a path fact of lesser cost C with the same value for X, Y (i.e. between the same pair of points). If there is such a fact, the costlier path fact is discarded. This aggregate selection is extremely important for efficiency — without it the program may run for ever, generating cyclic paths of increasing length. With this aggregate selection, along with the choice annotation @aggregate_selection $p(X, Y, P, C)(X, Y, C)any(P)$, a single source query on the program runs in time $O(E \cdot V)$, where there are E edge facts, and V nodes in the graph. As shown here, CORAL's aggregate selection mechanism can also be used to provide a version of the choice operator of LDL, but with altogether different semantics [18].

5.6 Inter-Module Calls

The interaction between modules merits some discussion. Suppose that p is a predicate defined in module M1, and p appears in the body of a rule of module M2. Evaluation within a rule proceeds left-to-right[5] and can be thought of as a nested-loops join. (While this is not entirely accurate with respect to pipelined evaluation, it is an accurate enough description for our purposes.) When evaluation reaches the p literal, a scan is opened on p. A p tuple retrieved by the scan is used to instantiate the rule. When evaluation returns to the p literal on backtracking[6], the scan on p is advanced to get the next p tuple.

This 'get-next-tuple' interface to a relation p via a scan is the only interface presented to M2 by any relation, regardless of the nature of the relation, or the evaluation technique used. For example, if p is a derived relation defined in another module, the interface is still the same as if p were a base relation.

The point at which the called module returns answers, depends on its evaluation mode. If the called module is pipelined, an answer is returned as soon as it is found, and the computation of the called module is suspended until another answer is requested by the caller. For materialized evaluation, the use of certain features, such as 'save module' and 'aggregate selections' can result in all answers being computed before any answers are returned by the called module. Otherwise, answers are returned at the end of each fixpoint iteration (in the called module) in which an answer is generated; further iterations are carried out if more answers are requested by the calling module.

6 Interface with C++

The CORAL system has been integrated with C++ in order to support a combination of imperative and declarative programming styles. We have extended C++ by providing a collection of new classes (relations, tuples, args and scan descriptors) along with associated functions. There is also a construct to embed CORAL commands in C++ code. This extended C++ can be used in conjunction with the declarative language features of CORAL in two distinct ways:

- Relations can be computed in a declarative style using declarative modules, and then manipulated in imperative fashion in extended C++ without breaking the relation abstraction. In this mode of usage, typically there is a main program written in C++ that calls upon CORAL for the evaluation of some relations defined in CORAL modules. The main program is compiled (after some preprocessing) and executed from the operating system command prompt; the CORAL interactive interface is not used.
- New predicates can be defined using extended C++. These predicates can be used in declarative CORAL code (and incrementally loaded).

Thus, declarative code can call extended C++ code and vice-versa. The above two modes are further discussed in the following sections.

6.1 Extensions to C++

CORAL provides a collection of classes and associated functions to programmers who use the (extended) C++ language. The new classes are:

[5] More generally, in an order determined by the optimizer or user.

[6] Backtracking is only intra-rule unless evaluation in M2 is pipelined.

Arg : All CORAL data types are subclasses of Arg. A number of functions are provided to convert between C++ primitive types and CORAL data types.

Tuple : A tuple is a list of args (i.e. arguments).

Relation : The Relation class allows access to relations from C++. Relation values can be constructed through a series of explicit inserts and deletes, or through a call to a declarative CORAL module. The associated methods allow manipulation of relation values from C++ without breaking the relation abstraction.

C_ScanDesc : This abstraction supports relational scans in C++ code. A C_ScanDesc object is essentially a cursor over a relation.

In addition to the new classes, any sequence of commands that can be typed in at the CORAL interactive command interface can be embedded in C++ code, bracketed by special delimiters. A file containing C++ code with embedded CORAL code must first be passed through a CORAL preprocessor and then compiled using a standard C++ compiler.

6.2 Defining New Predicates

As we have already seen, predicates exported from one CORAL module can be used freely in other modules. Sometimes, it may be desirable to define a predicate using extended C++, rather than the declarative language supported within CORAL modules. Extended C++ provides an export mechanism for this task.

The predicate definition can use all features of extended C++. The source file is pre-processed into a C++ file, and compiled to produce an object file. If this object file is consulted from the CORAL prompt, it is loaded into a newly allocated region in the data area of the executing CORAL system.[7]

7 Extensibility in CORAL

The implementation of the declarative language of CORAL is designed to be extensible. The user can define new abstract data types, new relation implementations, or new indexing methods, and use the query evaluation system with minimal changes. The user's program will, of course, have to be compiled and linked with the system code. We assume a set of standard operations on data types, and all abstract data types to be manipulated by CORAL must provide these operations.

7.1 Extensibility of Data Types

The type system in CORAL is designed to be extensible; the class mechanism and virtual functions provided by C++ help make extensibility clean and local. 'Locality' refers to the ability to extend the type system by adding new code, without modifying existing system code — the changes are thus local to the code that is added. All abstract data types should have certain virtual functions defined in their interface, and all system code that manipulates objects operates only via this interface. This ensures that the query evaluation system does not need to be modified or recompiled when a new abstract data type is defined. The required functions include a function for checking if two objects are equal, a function for printing the object, a function for re-creating objects from a printed representation, a function for hashing objects, and some memory management functions. A summary of the virtual functions that constitute the abstract data type interface is presented in [19]. In addition to creating the abstract data type, the user can define predicates to manipulate (and possibly display in novel ways) objects belonging to the abstract data types. These predicates must be registered with the system; registration is accomplished by a single command.

7.2 Extensibility of Access Structures

CORAL currently supports relations organized as linked lists, organized as hash tables, defined by rules, or defined by C++ functions. The interface code to relations makes no assumptions about the structure of relations, and is designed to make the task of adding new relation implementations easy. The 'get-next-tuple' interface between the query evaluation system and a relation is the basis for adding new relation implementations and index implementations in a clean fashion. The implementation of persistent relations using EXODUS illustrates the utility of such extensibility.

8 Related Systems

There are many similarities between CORAL and deductive database systems such as Aditi [28], EKS-V1 [29], LDL [13, 3], Glue-NAIL! [11, 15], Starburst SQL [12], DECLARE [9], ConceptBase [6] and LOLA [4]. However, there are several important differences, and CORAL extends all the above systems in the following ways:

1. CORAL supports a larger class of programs, including programs with non-ground facts, non-stratified negation and set-generation.
2. CORAL supports a wide range of evaluation techniques, and gives the user considerable control over the choice of techniques.
3. CORAL is extensible — new data and relation types and index implementations can be added without modifying the rest of the system.

EKS-V1 supports integrity constraint checking, hypothetical reasoning and provides some support for non-stratified aggregation [10]. ConceptBase supports DATALOG, along with locally stratified negation (but no set-generation), several object-oriented features, integrity constraint checking, and provides a one-way interface to C/Prolog, i.e. the imperative language can call ConceptBase, but not vice versa.

[7] That is, the new code is incrementally loaded into CORAL.

LOLA supports stratified programs, integrity constraints, several join strategies, and some support for type information. The host language of LOLA is Lisp, and it is linked to the TransBase relational database. Aditi gives primary importance to disk-resident data and supports several join strategies.

Unlike Glue-NAIL! and LDL, where modules have only a compile-time meaning and no run-time meaning, modules in CORAL have important run-time semantics, in that several run-time optimizations are done at the module level. Modules with run-time semantics are also found in several production rule systems (for example, RDL1 [8]). LDL++, a successor to LDL under development at MCC Austin, is reportedly also moving in the direction taken by CORAL in many respects. It will be partially interpreted, support abstract data types, and use a local semantics for choice (Carlo Zaniolo, personal communication). XSB is a system being developed at SUNY, Stony Brook. It will support several features similar to CORAL, such as non-ground terms and modularly stratified set grouping and negation. Program evaluation in XSB will use OLDTNF, which has been implemented by modifying the WAM (David S. Warren, personal communication). DECLARE and SDS are early efforts to commercialize deductive database technology.

In comparison to logic programming systems, such as various implementations of Prolog, CORAL provides better indexing facilities and support for persistent data. Most importantly, the declarative intended model semantics is supported (for all positive Horn clause programs, and a large class of programs with negation and aggregation as well).

9 Conclusions

The CORAL project is at a stage where one version of the system has been released in the public domain, and an enhanced version will soon be released. The effects of several design decisions are becoming increasingly evident. On the positive side, most of the decisions we made seem to have paid off with respect to simplicity and ease of efficient implementation.

Modular Design : The concept of modules in CORAL was in many ways the key to the successful implementation of the system. Given the ambitious goal of combining many evaluation strategies in an orthogonal fashion, the module mechanism appears to have been the ideal approach.

Annotations : The strategy of allowing the users to express control choices was a convenient approach to solving the otherwise difficult problem of making optimization decisions.

Extensibility : The decision to design an extensible system seems to have helped greatly in keeping our code clean and modular, in addition to its utility from an application development perspective.

System Architecture : The architecture concentrated on the design of a single user database system, leaving issues like transaction management, concurrency control and recovery to be handled by the EXODUS toolkit. Thus CORAL could build on these facilities that were already available, and focus instead on the subtleties of deductive databases and logic rules. The overall architecture was reasonably successful in breaking the problem of query processing into relatively orthogonal tasks.

On the negative side, some poor decisions were made, and some issues were not addressed adequately.

Type Information : CORAL makes no effort to use type information in its processing. No type checking or inferencing is performed at compile-time, and errors due to type mismatches lead to subtle run-time errors. Typing is a desirable feature, especially if the language is to be used to develop large applications. Type information can be used when compiling declarative programs, and can improve execution speed greatly. Typing is one of the issues addressed by a proposed extension to CORAL [25].

Memory Management : In an effort to make the system as efficient as possible for main-memory operations, copying of data has largely been replaced by sharing using pointers. This makes evaluation more efficient when using large data items such as lists, but it requires extensive memory management and garbage collection. For simplicity and uniformity, we used pointer-based sharing even for primitive data types such as integers, where the benefit is small, and the cost of memory management is large. The time and space cost of this approach has turned out to be rather high; hence we plan to modify the tuple structure to store data of primitive types in the tuple itself and thereby reduce memory management overheads. Another motivation for this change, related to EXODUS buffers, was discussed in Section 3.2.

There are a number of directions in which CORAL could be, and in some cases needs to be, extended. These include better support for persistent data, improved memory management, enhanced C++ interface features, object-oriented extensions and support for constraints. While performance measurements of a preliminary nature have been made, an extensive performance evaluation of CORAL is planned for the near future.

Acknowledgements

We would like to acknowledge our debt to Aditi, EKS-V1, LDL, NAIL!, SQL, Starburst, and various implementations of Prolog from which we have borrowed numerous ideas. We would also like to acknowledge the contributions of the following people. Per Bothner was largely responsible for the initial implementation of CORAL that served as the basis for subsequent development; Joseph Albert worked on some aspects of the set-manipulation code; Tarun Arora implemented several utilities and built-in libraries; Tom Ball implemented an early prototype of the seminaive evaluation system; Lai-chong Chan did the initial implementation of existential query optimization; Sumeer Goyal implemented embedded CORAL constructs in C++; Vish Karra implemented pipelining; Robert Netzer did the initial implementation of Magic rewriting; and Bill Roth implemented the explanation tool along with Tarun, and added user-interface features.

References

[1] I. Balbin and K. Ramamohanarao. A generalization of the differential approach to recursive query evaluation. *Journal of Logic Programming*, 4(3), September 1987.

[2] M. Carey, D. DeWitt, G. Graefe, D. Haight, J. Richardson, D. Schuh, E. Shekita, and S. Vandenberg. The EXODUS extensible DBMS project: An overview. In *Readings in Object-Oriented Databases*. Morgan-Kaufman, 1990.

[3] D. Chimenti, R. Gamboa, R. Krishnamurthy, S. Naqvi, S. Tsur, and C. Zaniolo. The LDL system prototype. *IEEE Transactions on Knowledge and Data Engineering*, 2(1):76–90, 1990.

[4] B. Freitag, H. Schűtz, and G. Specht. LOLA - a logic language for deductive databases and its implementation. In *Proceedings of 2nd International Symposium on Database Systems for Advanced Applications (DASFAA)*, 1991.

[5] E. Goto. Monocopy and associative algorithms in an extended lisp. Technical Report 74-03, Information Science Laboratory, Univ. of Tokyo, Tokyo, Japan, May 1974.

[6] M. Jeusfeld and M. Staudt. Query optimization in deductive object bases. In G. J.C. Freytag, G. Vossen and D. Maier, editors, *Query Processing for Advanced Database Applications*. Morgan-Kaufmann, 1993.

[7] D. Kemp, K. Ramamohanarao, and Z. Somogyi. Right-, left-, and multi-linear rule transformations that maintain context information. In *Proceedings of the International Conference on Very Large Databases*, pages 380–391, Brisbane, Australia, 1990.

[8] G. Kiernan, C. de Maindreville, and E. Simon. Making deductive database a practical technology: a step forward. In *Proceedings of the ACM SIGMOD Conference on Management of Data*, 1990.

[9] W. Kießling and H. Schmidt. DECLARE and SDS: Early efforts to commercialize deductive database technology. Submitted to the VLDB Journal., 1993.

[10] A. Lefebvre. Towards an efficient evaluation of recursive aggregates in deductive databases. In *Proceedings of the International Conference on Fifth Generation Computer Systems*, June 1992.

[11] K. Morris, J. D. Ullman, and A. Van Gelder. Design overview of the NAIL! system. In *Proceedings of the Third International Conference on Logic Programming*, 1986.

[12] I. S. Mumick, H. Pirahesh, and R. Ramakrishnan. Duplicates and aggregates in deductive databases. In *Proceedings of the Sixteenth International Conference on Very Large Databases*, Aug. 1990.

[13] S. Naqvi and S. Tsur. *A Logical Language for Data and Knowledge Bases*. Principles of Computer Science. Computer Science Press, New York, 1989.

[14] J. F. Naughton, R. Ramakrishnan, Y. Sagiv, and J. D. Ullman. Argument reduction through factoring. In *Proceedings of the Fifteenth International Conference on Very Large Databases*, pages 173–182, Amsterdam, The Netherlands, August 1989.

[15] G. Phipps, M. A. Derr, and K. A. Ross. Glue-NAIL!: A deductive database system. In *Proceedings of the ACM SIGMOD Conference on Management of Data*, pages 308–317, 1991.

[16] R. Ramakrishnan. Magic Templates: A spellbinding approach to logic programs. In *Proceedings of the International Conference on Logic Programming*, pages 140–159, Seattle, Washington, August 1988.

[17] R. Ramakrishnan, C. Beeri, and R. Krishnamurthy. Optimizing existential Datalog queries. In *Proceedings of the ACM Symposium on Principles of Database Systems*, pages 89–102, Austin, Texas, March 1988.

[18] R. Ramakrishnan, P. Bothner, D. Srivastava, and S. Sudarshan. CORAL: A database programming language. In J. Chomicki, editor, *Proceedings of the NACLP '90 Workshop on Deductive Databases*, October 1990. Available as Report TR-CS-90-14, Department of Computing and Information Sciences, Kansas State University.

[19] R. Ramakrishnan, P. Seshadri, D. Srivastava, and S. Sudarshan. The CORAL user manual: A tutorial introduction to CORAL. Manuscript, 1993.

[20] R. Ramakrishnan, D. Srivastava, and S. Sudarshan. Rule ordering in bottom-up fixpoint evaluation of logic programs. *IEEE Transactions on Knowledge and Data Engineering (to appear)*. A shorter version appeared in VLDB, 1990.

[21] R. Ramakrishnan, D. Srivastava, and S. Sudarshan. Controlling the search in bottom-up evaluation. In *Proceedings of the Joint International Conference and Symposium on Logic Programming*, 1992.

[22] R. Ramakrishnan, D. Srivastava, and S. Sudarshan. CORAL: Control, Relations and Logic. In *Proceedings of the International Conference on Very Large Databases*, 1992.

[23] R. Ramakrishnan, D. Srivastava, and S. Sudarshan. Efficient bottom-up evaluation of logic programs. In J. Vandewalle, editor, *The State of the Art in Computer Systems and Software Engineering*. Kluwer Academic Publishers, 1992.

[24] R. Ramakrishnan and S. Sudarshan. Top-Down vs. Bottom-Up Revisited. In *Proceedings of the International Logic Programming Symposium*, 1991.

[25] D. Srivastava, R. Ramakrishnan, P. Seshadri, and S. Sudarshan. CORAL++: Adding object-orientation to a logic database language. Submitted.

[26] S. Sudarshan and Raghu Ramakrishnan. Optimizations of bottom-up evaluation with non-ground terms. In *JICSLP'92 Post-Conference Workshop on Deductive Databases*, 1992.

[27] S. Tsur and C. Zaniolo. LDL: A logic-based data-language. In *Proceedings of the Twelfth International Conference on Very Large Data Bases*, pages 33–41, Kyoto, Japan, August 1986.

[28] J. Vaghani, K. Ramamohanarao, D. Kemp, Z. Somogyi, and P. Stuckey. The Aditi deductive database system. In *Proceedings of the NACLP'90 Workshop on Deductive Database Systems*, 1990.

[29] L. Vieille, P. Bayer, V. Küchenhoff, and A. Lefebvre. EKS-V1, a short overview. In *AAAI-90 Workshop on Knowledge Base Management Systems*, 1990.