

Greedy algorithms in Datalog

SERGIO GRECO

*Dipartimento Elettronica Informatica e Sistemistica, Università della Calabria,
87030 Rende, Italy
greco@si.deis.unical.it*

CARLO ZANIOLO

*Computer Science Department, University of California at Los Angeles,
Los Angeles, CA 90024, USA
zaniolo@cs.ucla.edu*

Abstract

In the design of algorithms, the greedy paradigm provides a powerful tool for solving efficiently classical computational problems, within the framework of procedural languages. However, expressing these algorithms within the declarative framework of logic-based languages has proven a difficult research challenge. In this paper, we extend the framework of Datalog-like languages to obtain simple declarative formulations for such problems, and propose effective implementation techniques to ensure computational complexities comparable to those of procedural formulations. These advances are achieved through the use of the choice construct, extended with preference annotations to effect the selection of alternative stable-models and nondeterministic fixpoints. We show that, with suitable storage structures, the differential fixpoint computation of our programs matches the complexity of procedural algorithms in classical search and optimization problems.

KEYWORDS: Datalog, greedy algorithms, query evaluation

1 Introduction

The problem of finding efficient implementations for declarative logic-based languages represents one of the most arduous and lasting research challenges in computer science. The interesting theoretical challenges posed by this problem are made more urgent by the fact that extrema and other non-monotonic constructs are needed to express many real-life applications, ranging from the ‘Bill of Materials’ to graph-computation algorithms.

Significant progress in this area has been achieved on the semantic front, where the introduction of the well-founded model semantics and stable-model semantics allows us to assign a formal meaning to most, if not all, programs of practical interest. Unfortunately, the computational problems remain largely unsolved: various approaches have been proposed to more effective computations of well-founded models and stable models (Van Gelder *et al.*, 1991; Gelfond & Lifschitz, 1988), but these fall far short of matching the efficiency of classical procedural solutions for say, algorithms that find shortest paths in graphs. In general, it is known that

determining whether a program has a stable model is NP-complete (Marek & Truszczyński, 1991).

Therefore, in this paper we propose a different approach: while, at the semantic level, we strictly adhere to the formal declarative semantics of logic programs with negation, we also allow the use of extended non-monotonic constructs with first order semantics to facilitate the task of programmers and compilers alike. This entails simple declarative formulations and nearly optimal executions for large classes of problems that are normally solved using greedy algorithms.

Greedy algorithms (Moret & Shapiro, 1993) are those that solve a class of optimization problems, using a control structure of a single loop, where at each iteration some element judged the 'best' at that stage is chosen and it is added to the solution. The simple loop hints that these problems are amenable to a fixpoint computation. The choice at each iteration calls attention to mechanisms by which nondeterministic choices can be expressed in logic programs. This framework also provides an opportunity of making, rather than blind choices, choices based on some heuristic criterion, such as greedily choosing the least (or most) among the values at hand when seeking the global minimization (or maximization) of the sum of such values. Following these hints, this paper introduces primitives for choice and greedy selection, and shows that classical greedy algorithms can be expressed using them. The paper also shows how to translate each program with such constructs to a program which contains only negation as nonmonotonic construct, and which defines the semantics of the original program. Finally, several classes of programs with such constructs are defined and it is shown that (i) they have stable model semantics, (ii) they are easily identifiable at compile time, and (iii) they can be optimized for efficient execution, i.e. they yield the same complexities as those expected from greedy algorithms in procedural programs. Thus, the approach provides a programmer with declarative tools to express greedy algorithms, frees him/her from many implementation details, yet guarantees good performance.

Previous work has shown that many non deterministic decision problem can be easily expressed using the nondeterministic construct *choice* in logic programs (Saccà & Zaniolo, 1990; Giannotti *et al.*, 1991). We showed (Giannotti *et al.*, 1999) that while the semantic of choice requires the use of negation under total stable model semantics, a stable model for these programs can be computed in polynomial time. In fact, choice in Datalog programs stratified with respect to negation achieves DB-Ptime completeness under genericity (Abiteboul *et al.*, 1994). In this paper, we further explore the ability of choice to express and support efficient computations, by specializing choice with optimization heuristics expressed by the *choice-least* and *choice-most* predicates. Then, we show that these two new built-in predicates enable us to express easily greedy algorithms; furthermore, by using appropriate data structures, the least-fixpoint computation of a program with choice-least and choice-most emulates the classical greedy algorithms, and achieves their asymptotic complexity.

A significant amount of excellent previous work has investigated the issue of how to express in logic and compute efficiently greedy algorithms, and, more in general, classical algorithms that require non-monotonic constructs. An incomplete

list include work by Sudarshan & Ramakrishnan (1991), Dietrich (1992), Van Gelder (1993) and Ganguly *et al.* (1995). This line of research was often motivated by the observation that many greedy algorithms can be viewed as optimized versions of transitive closures. Efficient computation of transitive closures is central to deductive database research, and the need for greedy algorithms is pervasive in deductive database applications and in more traditional database applications such as the Bill of Materials (Zaniolo *et al.*, 1997). In this paper, we introduce a treatment for greedy algorithms that is significant simpler and more robust than previous approaches (including that proposed in Greco *et al.* (1992)), where it was proposed to use the *choice* together with the built-in predicates *least* and *most*); it also treats all aspects of these algorithms, beginning from their intuitive formulation, and ending with their optimized expression and execution.

The paper is organized as follows. In section 2 we present basic definitions on the syntax and semantics of Datalog. In section 3, we introduce the notion of choice and the stable-model declarative semantics of choice programs. In section 4, we show how with this nondeterministic construct we can express in Datalog algorithms such as single-source reachability and Hamiltonian path. A fixpoint-based operational semantics for choice programs presented in section 5, and this semantics is then specialized with the introduction of the choice-least and choice-most construct to force greedy selections among alternative choices. In section 6, we show how the greedy refinement allow us to express greedy algorithms such as Prim's and Dijkstra's. Finally, in section 7, we turn to the implementation of choice, choice-least and choice-most programs, and show that using well-known deductive DB techniques, such as differential fixpoint, and suitable access structures, such as hash tables and priority queues, we achieve optimal complexity bounds for classical search problems.

2 Basic notions

In this section, we summarize the basic notions of Horn clauses logic, and its extensions to allow negative goals.

A *term* is a variable, a constant, or a complex term of the form $f(t_1, \dots, t_n)$, where t_1, \dots, t_n are terms. An *atom* is a formula of the language that is of the form $p(t_1, \dots, t_n)$ where p is a predicate symbol of arity n . A *literal* is either an atom (positive literal) or its negation (negative literal). A *rule* is a formula of the language of the form

$$Q \leftarrow Q_1, \dots, Q_m,$$

where Q is a atom (*head* of the rule) and Q_1, \dots, Q_m are literals (*body* of the rule). A term, atom, literal or rule is *ground* if it is variable free. A ground rule with empty body is a fact. A logic program is a set of rules. A rule without negative goals is called positive (a Horn clause); a program is called positive when all its rules are positive. A DATALOG program is a positive program not containing complex terms.

Let P be a program. Given two predicate symbols p and q in P , we say that p

directly depends on q , written $p < q$ if there exists a rule r in P such that p is the head predicate symbol of r and q occurs in the body of r . The binary graph representing this relation is called the *dependency graph* of P . The maximal strong components of this graph will be called *recursive cliques*. Predicates in the same recursive clique are *mutually recursive*. A rule is *recursive* if its head predicate symbol is mutually recursive with some predicate symbol occurring in the body.

Given a logic program P , the Herbrand universe of P , denoted H_P , is the set of all possible ground terms recursively constructed by taking constants and function symbols occurring in P . The Herbrand Base of P , denoted B_P , is the set of all possible ground atoms whose predicate symbols occur in P and whose arguments are elements from the Herbrand universe. A *ground instance* of a rule r in P is a rule obtained from r by replacing every variable X in r by a ground term in H_P . The set of ground instances of r is denoted by $ground(r)$; accordingly, $ground(P)$ denotes $\bigcup_{r \in P} ground(r)$. A (*Herbrand*) *interpretation* I of P is any subset of B_P . An *model* M of P is an interpretation that makes each ground instance of each rule in P *true* (where a positive ground atom is *true* if and only if it belongs to M and a negative ground atom is *true* if and only if it does not belong to M -total models). A rule in $ground(P)$ whose body is true w.r.t. an interpretation I will also be called *fireable* in I . Thus, a model for a program can be constructed by a procedure that starts from $I := \emptyset$ and adds to I the head of a rule $r \in ground(P)$ that is fireable in I (this operation will be called *firing* r) until no fireable rules remain. A model of P is *minimal* if none of its proper subsets is a model. Each positive logic program has a unique minimal model which defines its formal declarative semantics.

Given a program P and an interpretation M for P , we denote as $ground_M(P)$ the program obtained from $ground(P)$ by

1. removing every rule having as a goals some literal $\neg q$ with $q \in M$,
2. removing all negated goals from the remaining rules.

Since $ground_M(P)$ is a positive program, it has a unique minimal model. A model M of P is said to be *stable* when M is also the minimum model of $ground_M(P)$ (Gelfond & Lifschitz, 1988). A given program can have one or more stable (total) model, or possibly none. Positive programs, stratified programs (Apt *et al.*, 1988), locally stratified programs (Przymusinski, 1988) and weakly stratified programs (Przymusinska & Przymusinski, 1988) are among those that have exactly one stable model.

Let I be an interpretation for a program P . The *immediate consequence operator* $T_P(I)$ is defined as the set containing the heads of each rule $r \in ground(P)$ s.t. all positive goals of r are in I , and none of the negated goals of r , is in I .

3 Nondeterministic reasoning

Say that our university database contains a relation `student(Name, Major, Year)`, and a relation `professor(Name, Major)`. In fact, let us take a toy example that only has the following facts:

```

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 $\text{choice}((S), (P))$ must be added to force the selection of a unique advisor, out of the eligible advisors, for a student.

Example 3.1

Computation of unique advisors by choice rules

```

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

```

The computation of this rule gives for each student S a unique professor P .

The goal $\text{choice}((S), (P))$ 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).

The result of executing this rule is *nondeterministic*. It can either give a singleton relation containing the tuple ('JimBlack', ohm) or that containing the tuple ('JimBlack', bell).

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, $\text{foe}(P)$, called the *first order equivalent* of a choice program P . $\text{foe}(P)$ exhibits a multiplicity of stable models, each obeying the FDs defined by the choice goals. Each stable model for $\text{foe}(P)$ corresponds to an alternative set of answers for P and is called a *choice model* for P . $\text{foe}(P)$ is defined as follows:

Definition 3.1

(Saccà & Zaniolo, 1990) The first order equivalent version $\text{foe}(P)$ of a choice program P is obtained by the following transformation. Consider a choice rule r in P :

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

where

- (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, $\text{foe}(P)$ is constructed by transforming the original program P as follows:

1. Replace r with a rule r' obtained by substituting the choice goals with the atom $chosen_r(W)$:

$$r' : A \leftarrow B(Z), chosen_r(W),$$

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

$$chosen_r(W) \leftarrow B(Z), \neg diffchoice_r(W).$$

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

$$diffchoice_r(W) \leftarrow chosen_r(W'), Y_i \neq Y'_i,$$

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) $Y_i \neq Y'_i$ is true if $A \neq A'$, for some variable $A \in Y_i$ and its primed counterpart $A' \in Y'_i$.

The first order equivalent version of Example 3.1 is given in Example 3.2, which 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.

Example 3.2

The first order equivalent version of the rule in Example 3.1 is

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

In general, the program $foe(P)$ generated by the transformation discussed above has the following properties (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.

The stable models of $foe(P)$ are called *choice models* for P .

While the topic of operational semantics for choice Datalog programs will be further discussed in section 5, it is clear that choice programs can be implemented efficiently. Basically, the *chosen* atoms must be produced one-at-a-time and memoized in a table. The *diffchoice* atoms need not be computed and stored (*diffchoice* rules are not range restricted and their evaluation could produce huge results); rather, a goal $\neg diffchoice(t)$ can simply be checked dynamically against the table *chosen*. Since these are simple operations (actually quasi constant-time if an hash table is used), it follows that choice Datalog programs can be computed in polynomial time, and that rules with choice can be evaluated as efficiently as those without choice.

4 Computing with choice

Choice significantly extends the power of Datalog, and Datalog with stratified negation (Greco *et al.*, 1995; Giannotti *et al.*, 1991). In this paper we consider Datalog with the nondeterministic construct *choice*, although our framework can be easily extended to also consider stratified negation.

The following example presents a choice program that pairwise chains the elements of a relation $d(X)$, thus establishing a random total order on these elements.

Example 4.1

Linear sequencing of the elements of a set. The elements of the set are stored by means of facts of the form $d(Y)$.

```

succ(root,root).
succ(X,Y) ←      succ(.,X), d(Y),
                  choice((X),(Y)), choice((Y),(X)).

```

Here $\text{succ}(\text{root},\text{root})$ is the root of a chain linking all the elements of $d(Y)$. The transitive closure of succ thus defines a total order on the elements of d . Because of the ability of choice programs to order the elements of a set, Datalog with choice is P-time complete and can, for instance, express the parity query, i.e. determining if a relation has an even number of elements (Abiteboul *et al.*, 1994). This 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* (Chandra & Harel, 1982; Abiteboul *et al.*, 1994).

The expressive power of the choice construct has been studied in Giannotti *et al.* (1999) and Greco *et al.* (1995), where it is shown that it is more powerful than other nondeterministic constructs, including the witness operator (Abiteboul & Vianu, 1991), and the original version of choice proposed by Krishnamurthy & Naqui (1988), which is called static-choice, to distinguish it from the dynamic choice used here (Giannotti *et al.*, 1991). For instance, it has been shown (Giannotti *et al.*, 1991) that the task of ordering a domain or computing whether a relation contains an even number of elements (parity query) cannot be performed by positive programs with static choice or the witness operator (Abiteboul & Vianu, 1991).

In the rest of the paper, we will study nondeterministic queries combined with optimization criteria. For instance, our previous advisor example can be modified using optimized criteria to match students with candidate advisors. In the next example we present the general matching problem for bipartite graphs.

Example 4.2

Matching in a bipartite graph. We are given a bipartite graph $G = \langle (V_1, V_2), E \rangle$, i.e. a graph where nodes are partitioned into two subset V_1 and V_2 and each edge connect nodes in V_1 with nodes in V_2 . The problem consists to find a matching, i.e., a subset E' of E such that each node in V_1 is joined with at most one node in V_2 and *vice versa*.

```

matching(X,Y) ←  g(X,Y,C), choice((Y),(X)),
                  choice((X),(Y)), choice((X),(C)).

```

Here a fact $g(x, y, c)$ denotes that there is an edge with cost c joining the node $x \in V_1$ with the node $y \in V_2$.

In section 6, we will consider the related optimization problem, of finding a matching such that the sum of all Cs is minimized or maximized¹.

Example 4.3

Rooted spanning tree. We are given an undirected graph where an edge joining two nodes, say x and y , is represented by means of two facts $g(x, y, c)$ and $g(y, x, c)$, where c is the cost. A spanning tree in the graph, starting from the source node a , can be expressed by means of the following program:

```

st(root, a, 0).
st(X, Y, C) ←  st(., X, .), g(X, Y, C), Y ≠ a, Y ≠ X,
                choice((Y), (X)), choice((Y), (C)).

```

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

```

g(a, b, 1).      g(b, a, 1).
g(b, c, 2).      g(c, b, 2).
g(a, c, 3).      g(c, a, 3).

```

After the exit rule adds $st(\text{root}, a, 0)$, the recursive rule could add $st(a, b, 1)$ and $st(a, c, 3)$ along with the two tuples $\text{chosen}(a, b, 1)$ and $\text{chosen}(a, c, 3)$ in the chosen table. No further arc can be added after those, since the addition of $st(b, c, 2)$ or $st(c, b, 2)$ would violate the FD that follows from $\text{choice}((Y), (X))$ enforced through the chosen table. However, since $st(\text{root}, a, 0)$, 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 $st(c, a, 3)$ or $st(b, a, 1)$, 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: $M_1 = \{st(a, b, 1), st(b, c, 2)\} \cup G$, $M_2 = \{st(a, b, 1), st(a, c, 3)\} \cup G$ and $M_3 = \{st(a, c, 3), st(c, b, 2)\} \cup G$.

Example 4.4

Single-Source Reachability. Given a direct graph where the arcs are stored by means of tuples of the form $g(x, y, c)$, the set of nodes reachable from a node a can be defined by the following program:

```

reach(a, 0).
reach(Y, C) ←  reach(X, C1), g(X, Y, C2), Y ≠ a,
                C = C1 + C2, choice((Y), (C)).

```

¹ Given that the pair $X \rightarrow Y, X \rightarrow C$ is equivalent to $X \rightarrow Y, C$, the last rule in the previous example can also be written as follows:

```

matching(X, Y) ← g(X, Y, C), choice((Y), (X)), choice((X), (Y, C)).

```


Once the cost arguments are eliminated from these rules, we obtain the usual transitive-closure-like program, for which the fixpoint computation terminates once all nodes reachable from node a are found, even if the graph contains cycles. However, if the choice goal were eliminated, the program of Example 4.4 could become nonterminating on a cyclic graph.

In the next example, we have a complete undirected labeled graph G , represented by facts $g(x, y, c)$, where the label c typically represents the cost of the edge. A simple path is a path passing through a node at most once. A Hamiltonian path is a simple path reaching each node in the graph. Then, a simple path can be constructed as follows:

Example 4.5

The simple path problem. When the arc from X to Y is selected, we must make sure that the ending node Y had not been selected and the starting node X is connected to some selected node. The choice constraints, and the goals $\text{spath}(\text{root}, Z, 0), Y \neq Z$ to avoid returning to the initial node, ensure that a simple path is obtained.

$$\begin{aligned} \text{spath}(\text{root}, X, 0) &\leftarrow g(X, _ , _), \text{choice}(_ , X). \\ \text{spath}(X, Y, C) &\leftarrow \text{spath}(_ , X, _), g(X, Y, C), \text{spath}(\text{root}, Z, 0), Y \neq Z, \\ &\quad \text{choice}(X, (Y)), \text{choice}(Y, (X)), \text{choice}(Y, (C)). \end{aligned}$$

When G is a complete graph, the simple path produced by this program is Hamiltonian (i.e. touches all the nodes). In many applications, we need to find a minimum-cost Hamiltonian path; this is the Traveling Salesman Problem (TSP) discussed in section 6.

The next program presents a problem consisting in the selection of a set of elements satisfying a constraint. The optimized version of this problem is the well-known knapsack problem.

5 Fixpoint semantics

5.1 Choice programs

Let I be an interpretation for a program P ; the *immediate consequence operator* $T_P(I)$ is defined as the set containing the heads of each rule $r \in \text{ground}(P)$ s.t. all positive goals of r are in I , and none of the negated goals of r , is in I . For a choice program P , with first order equivalent $\text{foe}(P)$, let us denote by T_{P_c} the immediate consequence operator associated with the rules defining the predicate chosen in $\text{foe}(P)$ (these are the rules with the $\neg\text{diffchoice}$ goals) and let T_{P_d} denote the immediate consequence for all the other rules in $\text{foe}(P)$ (for positive choice programs these are Horn clauses).

Therefore, we have that, for any interpretation I of $\text{foe}(P)$:

$$T_{\text{foe}(P)}(I) = T_{P_d}(I) \cup T_{P_c}(I).$$

Following Giannotti *et al.* (1999), we can now introduce a general operator for computing the nondeterministic fixpoints of a choice program P . We will denote by FD_P the functional dependencies defined by the choice goals in P .

Definition 5.1

Given a choice program P , its *nondeterministic immediate consequence operator* Ψ_P is a mapping from an interpretation of $foe(P)$ to a set of interpretations of $foe(P)$ defined as follows:

$$\Psi_P(I) = \{ T_{P_D}^{\uparrow\omega}(I \cup \Delta C) \cup \Delta C \mid \Delta C \in \Gamma_P(I) \} \quad (1)$$

where $\Gamma_P(I) = \{\emptyset\}$ if $T_{P_C}(I) = \emptyset$, and otherwise:

$$\Gamma_P(I) = \{ \Delta C \mid \emptyset \subset \Delta C \subseteq T_{P_C}(I) \setminus I \text{ and } I \cup \Delta C \models FD_P \} \quad (2)$$

with $I \cup \Delta C \models FD_P$ denoting that $I \cup \Delta C$ satisfies the dependencies in FD_P .

Therefore, the Ψ_P operator is basically the composition of two operators. Given an interpretation I , the first operator computes all the admissible subsets of $\Delta C \subseteq T_{P_C}(I)$, i.e. those where $I \cup \Delta C$ obeys the given FDs; the second operator derives the logical consequence for each admissible subset using the ω -power of T_{P_D} .

The definition of $\Gamma_P(I)$ is such that ΔC is not empty iff $T_{P_C}(I) \setminus I$ is not empty; thus, if there are possible new choices, then at least one has to be taken. The Ψ_P operator formalizes a single step of a bottom-up computation of a choice program. Instead of defining the powers of Ψ_P , it is technically more convenient to define directly the notion of a nondeterministic computation based on the Ψ_P operator.

Observe that given the presence of the constraint, $I \cup \Delta C \models FD_P$, we can eliminate the $\neg\text{diffchoice}$ goal from the chosen rules. In fact, if $T_{P'_C}$ denotes the immediate consequence operator for the chosen rules without the $\neg\text{diffchoice}$ goals, then $T_{P'_C}$ can replace T_{P_C} in equation (2).

Definition 5.2

Given a choice program P , an *inflationary choice fixpoint computation* for P , is a sequence $\langle I_n \rangle_{n \geq 0}$ of interpretations such that:

- (i) $I_0 = \emptyset$,
- (ii) $I_{n+1} \in \Psi_P(I_n)$, for $n \geq 0$.

Inasmuch as every sequence $\langle I_n \rangle_{n \geq 0}$ is monotonic, it has a unique limit for $n \rightarrow \infty$; this limit will be called an *inflationary choice fixpoint* for the choice program P . Thus, we have the following result:

Theorem 5.1 ((Giannotti et al., 1991))

Let P be a Datalog program with choice, and M a Herbrand interpretation for $foe(P)$. Then M is a choice model for P iff M is an inflationary choice fixpoint for P .

Moreover, the inflationary choice fixpoint is *sound* (every result is a choice model) and *complete* (for each choice model there is some inflationary choice fixpoint computation producing it). For logic programs with infinite Herbrand universe, an additional assumption of *fairness* is needed to ensure completeness (Giannotti *et al.*, 1999). As customary for database queries, computational complexity is evaluated with respect to the size of the database. Then, we have the following result:

Theorem 5.2 ((Giannotti et al., 1991))

Let P be a choice Datalog program. Then, the data complexity of computing a choice model for P is polynomial time.

Therefore, for a choice Datalog program, P , the computation of one of the stable models for $foe(P)$ can be performed in polynomial time using the Choice Fixpoint Computation. This contrasts with the general intractability of finding stable models for general programs: in fact, we know that checking if a Datalog program with negation has a stable model is NP-complete (Marek & Truszczyński, 1991).

Therefore, the choice construct allows us to capture a special subclass of programs that have a stable model semantics but are amenable to efficient implementation and are appealing to intuition. Implementing these programs only requires memorization of the *chosen* predicates; from these, the *diffchoice* predicates can be generated on-the-fly, thus eliminating the need to store *diffchoice* explicitly. Moreover, the model of memorizing tables to enforce functional dependencies provides a simple enough metaphor for a programmer to make effective usage of this construct without having to become cognizant on the subtleties of non-monotonic semantics. We conclude by mentioning that, although we are considering (positive) choice Datalog programs, our framework can be trivially extended to also consider stratified negation. The computation of a choice model for a stratified choice program can be carried out by partitioning the program into an ordered number of suitable subprograms (called 'strata') and computing the choice fixpoints of every stratum in their order.

5.2 Greedy choice

Definition 5.1 leaves quite a bit of latitude in the computation of Δ (equation (2)). This freedom can be used to select Δ s that have additional properties. In particular, we want to explore specializations of this concept that trade nondeterministic completeness (which is only of abstract interest to a programmer) in return for very concrete benefits, such as expressive power and performance. For instance, in the specialization called *Eager Choice* (Giannotti et al., 1991), a maximal ΔC is used in equation (2). This results in a significant increase in expressive power, as demonstrated by the fact that negation can be emulated by eager choice (Giannotti et al., 1991, 1999).

In this paper, we focus on a specialization of choice called *greedy choice*; our interest in this constructs follows from the observation that it is frequently desirable to select a value that is the *least* (or the *most*) among the possible values and still satisfy the FDs defined by the choice atoms.

A choice-least (resp. choice-most) atom is of the form `choice-least((X), (C))` (resp. `choice-most((X), (C))`) where X is a list of variables and C is a single variable ranging over an ordered domain. A rule may have at most one choice-least or one choice-most atom. A goal `choice-least((X), (C))` (resp. `choice-most((X), (C))`) in a rule r can be used to denote that the FD defined by the atom `choice((X), (C))` is to be satisfied – the declarative semantics of choice, choice-least and choice-most coincide. For instance, a rule of the form

$$p(X, Y, C) \leftarrow q(X, Y, C), \text{choice}((X), (Y)), \text{choice-least}((X), (C)).$$

defines the FD $X \rightarrow Y, C$ on the possible instances of p . Thus, assuming that q is defined by the facts $q(a, b, 1)$ and $q(a, c, 2)$, from the above rule we can derive either $p(a, b, 1)$ or $p(a, c, 2)$. Moreover, the choice-least goal introduces some heuristic in the computation to derive only $p(a, b, 1)$. This means that, by using choice-least and choice-most predicates, we introduce some preference criteria on the stable models of the program. The ‘greedy’ fixpoint computation permit us to compute a ‘preferred’ stable model.

We can now define a *choice-least rule* (resp. *choice-most rule*) as one that contains one choice-least (resp. one choice-most) goal, and zero or more choice goals. Moreover, we also assume that our programs contain either choice-least or choice-most rules. A program that contains choice-least rules (choice-most rules) and possibly other rules with zero or more choice goals is called a *choice-least* program (a *choice-most* program). Choice-least and choice-most programs have dual properties; thus in the rest of the paper we will often mention the properties of one kind of program with the understanding that the corresponding properties of the other are implicitly defined by this duality.

The correct computation of *choice-least* programs can be thus defined by specializing the nondeterministic immediate consequence operator by (i) ensuring that Δ is a singleton set, containing only one element (ii) ensuring that a least-cost tuple among those that are candidates is chosen.

Formally, we can use as our starting point the *lazy version of choice* where Δ is specialized into a singleton set δ . The specialized version of Ψ_P so derived will be denoted Ψ_P^{lazy} ; as proven in Giannotti *et al.* (1991), the inflationary choice fixpoint restricted using Ψ_P^{lazy} operators still provides a sound and nondeterministically complete computation for the choice models of P .

We begin by decomposing Ψ_P^{lazy} in three steps:

Definition 5.3

Lazy Immediate-Consequence Operator (LICO).

Let P be a choice program and I an interpretation of P . Then $\Psi_P(I)$ for P is defined as follows:

$$\begin{aligned} \Theta_I &= \{\delta \in T_{P_c}(I) \setminus I \mid I \cup \{\delta\} \models FD_P\} \\ \Gamma_P^{lazy}(I) &= \{I \cup \{\delta\} \mid \delta \in \Theta_I\} \cup \{I \mid \Theta_I = \emptyset\} \\ \Psi_P^{lazy}(I) &= \{T_{P_d}^{\uparrow\omega}(J) \mid J \in \Gamma_P^{lazy}(I)\}. \end{aligned}$$

Given an interpretation I , a set $\Delta \in \Gamma_P(I)$ and two tuples $t_1, t_2 \in \Delta$. We say that $t_1 < t_2$ if both tuples are inferred only by choice-least rules and the cost of t_1 is lesser than the cost of t_2 . Further, we denote with *least*(Δ) the set of tuples of Δ with least cost, i.e. $\text{least}(\Delta) = \{t \mid t \in \Delta \text{ and } \nexists u \in \Delta \text{ s.t. } u < t\}$.

Therefore, the implementation of greedy algorithms follows directly from replacing $\delta \in \Theta_I$ with $\delta \in \text{least}(\Theta_I)$.

Definition 5.4

Least-Cost Immediate-Consequence Operator.

Let P be a choice program and I an interpretation of P . Then $\Psi_P^{least}(I)$ for P is defined as follows:

$$\begin{aligned} \Theta_I &= \{ \delta \in T_{PC}(I) \setminus I \mid I \cup \{ \delta \} \models FD_P \} \\ \Gamma_P^{least}(I) &= \{ I \cup \{ \delta \} \mid \delta \in least(\Theta_I) \} \cup \{ I \mid \Theta_I = \emptyset \} \\ \Psi_P^{least}(I) &= \{ T_{P_D}^{\uparrow \omega}(J) \mid J \in \Gamma_P^{least}(I) \}. \end{aligned}$$

Ψ_P^{least} will be called the *Least-Cost Immediate-Consequence Operator*.

Likewise, we have the dual definition of the *Most-Cost Immediate-Consequence Operator*.

Definition 5.5

Let P be a program with choice and choice-least goals. An *inflationary least choice fixpoint computation* (LFC) for P , is a sequence $\langle I_n \rangle_{n \geq 0}$ of interpretations such that:

- (i) $I_0 = \emptyset$,
- (ii) $I_{n+1} \in \Psi_P^{least}(I_n)$, for $n \geq 0$.

Thus, all the tuples that do not violate the given FDs (including the FDs implied by least) are considered, and one is chosen that has the least value for the cost argument.

Theorem 5.3

Let P be a Datalog program with choice and choice.least. Then,

1. Every inflationary least choice fixpoint for P is a choice model for P .
2. Every inflationary least choice fixpoint of P can be computed in polynomial time.

Proof

For the first property, observe that every computation of the inflationary least choice fixpoint is also a computation of the lazy choice fixpoint. Therefore every inflationary least choice fixpoint for P is a choice model for P .

The second property follows from the fact that the complexity of the inflationary lazy choice fixpoint is polynomial time. Moreover, the cost of selecting a tuple with least cost is also polynomial. Therefore, the complexity of inflationary least choice fixpoint is also polynomial. \square

While the inflationary choice fixpoint computation is sound and complete with respect to the declarative stable-model semantics the inflationary least (most) choice fixpoint computation is sound but no longer complete; thus there are choice models that are never produced by this computation. Indeed, rather than following a ‘don’t care’ policy when choosing among stable models, we make greedy selections between the available alternatives. For many problems of interest, this greedy policy is sufficient to ensure that the resulting models have some important optimality

properties, such as the minimality of the sum of cost of the edges. The model so constructed, will be called *greedy choice models*².

6 Greedy algorithms

In a system that adopts a concrete semantics based on *least choice fixpoint*, a programmer will specify a `choice-least((X), (Y))` goal to ensure that only particular choice models rather than arbitrary ones are produced, through the greedy selection of the least values of Y at each step. Thus, an optimal matching in a directed graph problem can be expressed as follows:

Example 6.1

Optimal matching in a bipartite graph

```
opt_matching(X, Y) ← g(X, Y, C), choice((Y), (X)),
                    choice((X), (Y)), choice-least((X), (C)).
```

Observe that this program is basically that of Example 4.2 after that the choice goal with a cost argument has been specialized to a choice-least goal.

The specialization of choice goals into *choice-least* or *choice-most* goals yields a convenient and efficient formulation of many greedy algorithms, such as Dijkstra's shortest path and Prim's minimum-spanning tree algorithms discussed next.

The algorithm for finding the minimum spanning tree in a weighted graph, starting from a source node a, can be derived from the program of Example 4.3 by simply replacing the goal `choice((Y), (C))` with `choice-least((Y), (C))` yielding the well-known Prim's algorithm.

Example 6.2

Prim's Algorithm

```
st(root, a, 0).
st(X, Y, C) ← st(., X, .), g(X, Y, C), Y ≠ a,
              choice((Y), (X)), choice-least((Y), (C)).
```

Analogously, the algorithm for finding the shortest path in a weighted digraph, starting from a source node a, can be derived from the program of Example 4.4 by simply replacing the goal `choice((Y), (C))` with `choice-least((Y), (C))`, yielding the well-known Dijkstra's algorithm, below.

Example 6.3

Dijkstra's algorithm

```
dj(a, 0).
dj(Y, C) ← dj(X, C1), g(X, Y, C2), Y ≠ a,
           C = C1 + C2, choice-least((Y), (C)).
```

² In terms of relation between declarative and operational semantics, the situation is similar to that of pure Prolog programs, where the declarative semantics is defined by the set of all legal SLD-trees, but then one particular tree will be generated instead of others according to some preference criterion.

Consider now the program of Example 4.1, which chains the elements of a domain $d(X)$ in an arbitrary order. Say now that a particular lexicographical order is pre-defined and we would like to sort the elements of $d(X)$ accordingly. Then, we can write the rules as follows:

Example 6.4

Sequencing the elements of a relation in decreasing order

```

succ(root, root).
succ(X, Y) ←      succ(−, X), d(Y),
                   choice-most((X), (Y)), choice((Y), (X)).

```

Greedy algorithms often provide efficient approximate solutions to NP-complete problems; the following algorithm yields heuristically effective approximations of optimal solutions for the traveling salesperson problem (Papadimitriou & Steiglitz, 1975).

Example 6.5

Greedy TSP. Given a complete undirected graph, the exit rule simply selects an arbitrary node X , from which to start the search. Then, the recursive rule greedily chooses at each step an arc (X, Y, C) of least cost C having X as its end node.

```

spath(root, X, 0) ← node(X), choice(( ), X).
spath(X, Y, C) ←   spath(−, X, −), g(X, Y, C),
                   spath(root, Z, 0), Y ≠ Z,
                   choice((X), (Y)), choice((Y), (X)),
                   choice-least((Y), (C)).

```

Observe that the program of Example 6.5 was obtained from that of Example 4.5 by replacing a choice goal with its choice-least counterpart.

Example 6.6

While we have here concentrated on graph optimization problems, greedy algorithms are useful in a variety of other problems. For instance, in Greco & Zaniolo (1998), a greedy solution to the well-known knapsack problem is presented, consisting of finding a set of items whose total weight is lesser than a given value (say 100) and whose cost is maximum. This is an NP-complete problem and, therefore, the optimal solution requires an exponential time (assuming $P \neq NP$) but an approximate solution carried out by means of a greedy computation, which selects at each step the item with maximum *value/weight* ratio.

In conclusion, we have obtained a framework for deriving and expressing greedy algorithms (such as Prim's algorithm) characterized by conceptual simplicity, logic-based semantics, and short and efficient programs; we can next turn to the efficient implementation problem for our programs.

7 Implementation and complexity

A most interesting aspect of the programs discussed in this paper is that their stable models can be computed very efficiently. In the previous sections, we have seen that the exponential intractability of stable models is not an issue here: our greedy fixpoint computations are always polynomial-time in the size of the database. In this section, we show that the same asymptotic complexity obtainable by expressing the algorithms in procedural languages can be obtained by using comparable data structures and taking advantage of syntactic structure of the program.

In general, the computation consists of two phases: compilation and execution. All compilation algorithms discussed here execute with time complexity that is polynomial in the size of the programs. Moreover, we will assume, as it is customarily done (Zaniolo *et al.*, 1997), that the size of the database dominates that of the program. Thus, execution costs dominate the compilation costs, which can thus be disregarded in the derivation of the worst case complexities.

We will use compilation techniques, such as the differential fixpoint computation, that are of common usage in deductive database systems (Zaniolo *et al.*, 1997). Also, we will employ suitable storage structures, such as hash tables to support search on keys, and priority queues to support choice-least and choice-most goals.

We assume that our programs consist of a set of mutually recursive predicates. General programs can be partitioned into a set of subprograms where rules in every subprogram defines a set of mutually recursive predicates. Then, subprograms are computed according to the topological order defined by the dependencies among predicates, where tuples derived from the computation of a subprogram are used as database facts in the computation of the subprograms that follow in the topological order.

7.1 Implementation of programs with choice

Basically, the *chosen* atoms need to be memorized in a set of tables $chosen_r$ (one for each $chosen_r$ predicate). The *diffchoice* atoms need not be computed and stored; rather, a goal $\neg diffchoice_r(\dots)$ can simply be checked dynamically against the table $chosen_r$. We now present how programs with choice can be evaluated by means of an example.

Example 7.1

Consider Example 6.4 again:

$$\begin{aligned} s_1 &: \text{succ}(\text{root}, \text{root}). \\ s_2 &: \text{succ}(X, Y) \leftarrow \text{succ}(_, X), d(Y), \\ &\quad \text{choice-most}((X), (Y)), \text{choice}((Y), (X)). \end{aligned}$$

According to our definitions, these rules are implemented as follows:

$r_1 : \text{succ}(\text{root}, \text{root}).$
 $r_2 : \text{succ}(X, Y) \leftarrow \text{succ}(_, X), \text{d}(Y), \text{chosen}(X, Y).$
 $r_3 : \text{chosen}(X, Y) \leftarrow \text{succ}(_, X), \text{g}(X, Y, C), \neg \text{diffchoice}(X, Y).$
 $r_4 : \text{diffchoice}(X, Y) \leftarrow \text{chosen}(X, Y'), Y' \neq Y.$
 $r_5 : \text{diffchoice}(X, Y) \leftarrow \text{chosen}(X', Y), X' \neq X.$

(Strictly speaking, the *chosen* and *diffchoice* predicates should have been added the subscript s_2 for unique identification. However, we dispensed with that, since there is only one choice rule in the source program and no ambiguity can occur.) The *diffchoice* rules are used to enforce the functional dependencies $X \rightarrow Y$ and $Y \rightarrow X$ on the chosen tuples. These conditions can be enforced directly from the stored table $\text{chosen}(X, Y)$ by enforcing the following constraints ³:

$\leftarrow \text{chosen}(X, Y), \text{chosen}(X, Y'), Y' \neq Y$
 $\leftarrow \text{chosen}(X, Y), \text{chosen}(X', Y), X' \neq X$

that are equivalent to the two rules defining the predicate $\neg \text{diffchoice}$. Thus, rules r_4 and r_5 are never executed directly, nor is any *diffchoice* atom ever generated or stored. Thus, we can simply eliminate the *diffchoice* rules in the computation of our program $\text{foe}(P) = P_C \cup P_D$. In addition, as previously observed, we can eliminate the goal $\neg \text{diffchoice}$ from the chosen rules without changing the definition of *LICO* (the Lazy Immediate-Consequence Operator introduced in Definition 5.3). Therefore, let P'_D denote P_D after the elimination of the *diffchoice* rules, and let P'_C denoted the rules in P_C after the elimination of their negated *diffchoice* goals; then, we can express our LICO computation as follows:

$$\begin{aligned}
 \Theta_I &= \{ \delta \in T_{P'_C}(I) \setminus I \mid I \cup \{ \delta \} \models FD_P \} \\
 \Gamma_P^{\text{lazy}}(I) &= \{ I \cup \{ \delta \} \mid \delta \in \Theta_I \} \cup \{ I \mid \Theta_I = \emptyset \} \\
 \Psi_P^{\text{lazy}}(I) &= \{ T_{P'_D}^{\uparrow \omega}(J) \mid J \in \Gamma_P(I) \}.
 \end{aligned}$$

Various simplifications can be made to this formula. For the program of Example 7.1, P'_D consists of the exit rule r_1 , which only needs to be fired once, and of the rule r_2 , where the variables in choice goals are the same as those contained in the head. In this situation, the head predicate and the *chosen* predicate can be stored in the same table and $T_{P'_D}$ is implemented at no additional cost as part of the computation of *chosen*.

Consider now the implementation of a table chosen_x . The keys for this table are the left sides of the choice goals: X and Y for the example at hand. The data structures needed to support search and insertion on keys are well known. For main memory, we can use hash tables, where searching for a key value, and inserting or deleting an entry can be considered constant-time operations. Chosen tuples are stored into a table which can be accessed by means of a set of hash indexes. More specifically, for each functional dependency $X \rightarrow Y$ there is an hash index on the attributed specified by the variables in X .

³ A constraint is a rule with empty head which is satisfied only if its body is false.

7.2 Naive and semi-naive implementations

For Example 7.1, the application of the LICO to the empty set, yields $\Theta_{I_0} = \emptyset$; then, from the evaluation of the standard rules we get the set $\Psi_P^{lazy}(\emptyset) = \{p(\text{nil}, a)\}$. At the next iteration, we compute Θ_{I_1} and obtain all arcs leaving from node *a*. One of these arcs is chosen and the others are discarded, as it should be since they would otherwise violate the constraint $X \leftrightarrow Y$. This naive implementation of Ψ generates no redundant computation for Example 7.1; similar considerations also hold for the simple path program of Example 4.5. In many situations, however, tuples of Θ_I computed in one iteration, also belong to Θ_I in the next iteration, and memorization is less expensive than recomputation. Symbolic differentiation techniques similar to those used in the seminaive fixpoint computation, can be used to implement this improvement (Zaniolo *et al.*, 1997), as described below.

We consider the general case, where a program can have more than one mutually recursive choice rule and we need to use separate `chosenr` tables for each such rule. For each choice rule *r*, we also store a table `thetar` with the same attributes as `chosenr`. In `thetar`, we keep the tuples which are future candidates for the table `chosenr`.

We update incrementally the content of the tables `thetar` as they were concrete views, using differential techniques. In fact, $\Theta_r = \theta_r \sqcup \text{theta}_r$, where `thetar` is the table accumulation for the ‘old’ Θ_r tuples and θ_r is the set of ‘new’ Θ_r tuples generated using the differential fixpoint techniques. Finally, Θ_I in the LICO is basically the union of the Θ_r for the various choice rules *r*.

With P'_C be the set of chosen rules in $foe(P)$, with the `¬diffchoice` goal removed; let T_r denote the immediate consequence operator for a rule $r \in P'_C$. Also, P'_D will denote $foe(P)$ after the removal of the chosen rules and of the `diffchoice` rules: thus P'_D is P_D without the `diffchoice` rules.

The computation of Ψ_P^{lazy} can then be expressed by means of the algorithm reported in Figure 1.

In fact, the basic computation performed by our algorithm is operational translation of Ψ_P^{lazy} , enhanced with the differential computation of Θ_r . At Step 0, the non-choice rules are computed starting from the empty set. This corresponds to the computation of the non-recursive rules (exit rules) in all our examples, but Examples 4.2 and 4.5 for which the exit rules are choice rules and are first computed at Step 1.

In Step 1 and Step 3, of this algorithm, we used the function `conflictr(S,R)` defined next. Let *S* and *R* be two union-compatible relations, whose attribute sets contain the left sides of the choice goals in *r*, i.e., the unique keys of `chosenr` (*X* and *Y* for the example at hand). Then, `conflictr(S,R)` is the set of tuples in *S* whose `chosenr`-key values are also contained in *R*.

Now, Step 1 brings up to date the content of the `thetar` table, while ensuring that this does not contain any tuple conflicting with tuples in `chosenr`.

Symbolic differentiation techniques are used to improve the computation of $T_r(I) \setminus \text{theta}_r$ in recursive rules at Step 1 (ii) (Zaniolo *et al.*, 1997). This technique is particularly simple to apply to a recursive linear rule where the symbolic differentiation yields the same rule using, instead of the tuples of the whole predicate,

Algorithm 7.1

Semi-naive computation of a choice model.

Input: Choice program P .**Output:** Choice model I for $foe(P)$.**begin**

0 Initialization.

For every $r \in P'_C$ set $\text{chosen}_r = \text{theta}_r = \emptyset$;Set: $I := T_{P'_D}^{\uparrow\omega}(\emptyset)$;1 **Repeat**(i) Select an unmarked arbitrary $r \in P'_C$ and mark r (ii) Compute: $\theta_r = (T_r(I) \setminus \text{theta}_r) \setminus \text{conflict}(T_r(I) \setminus \text{theta}_r, \text{chosen}_r)$;(iii) Add θ_r to theta_r **Until** $\text{theta}_r \neq \emptyset$ or all rules in P'_C are marked;2 **If** $\text{theta}_r = \emptyset$ **Return** I ;3 (i) Select an arbitrary $x \in \text{theta}_r$, and move x from theta_r to chosen_r ;(ii) With $\delta = \{x\}$, delete from the selected table theta_r every tuple in $\text{conflict}(\text{theta}_r, \delta)$.4 Set: $I = T_{P'_D}^{\uparrow\omega}(I \cup \delta)$, then unmark all P'_C rules and resume from Step 1.**end.**

Fig. 1. Semi-naive computation.

the delta-tuples computed in the last step. All our examples but Examples 4.5 and 6.5 involve linear rule. The quadratic choice rule in Example 4.5 is differentiated into a pair of rules. In all examples, the delta-tuples are as follows:

- (i) The tuples produced by the exit rules at Step 0
- (ii) The new tuples produced at Step 4 of the last iteration. For all our examples, Step 4 is a trivial step where $T_{P'_D}^{\uparrow\omega}(I \cup \delta) = I \cup \delta$; thus, δ is the new value produced at Step 4.

Moreover, if r corresponds to a non-recursive, as the first rule in Examples 4.2 and 4.5, then this is only executed once with $\text{theta}_r = \emptyset$.

At Step 2, we check the termination condition, $\Theta_I = \emptyset$, i.e., $\Theta_r = \emptyset$ for all r in P_C .

Step 3 (ii) eliminates from theta_r all tuples that conflict with the tuple δ (including the tuple itself).

Therefore, Algorithm 7.1 computes a stable model for $foe(P)$ since it implements a differential version of the operator Ψ_P^{lazy} , and applies this operator until saturation.

We can now compute the complexity of the example programs of section 4. For graphs, we denote by n and e , respectively, the number of their nodes and edges.

Complexity of rooted spanning-tree algorithm: Example 4.3

The number of chosen tuples is bounded by $O(n)$, while the number of tuples computed by the evaluation of body rules is bounded by $O(e)$, since all arcs connected to the source node are visited exactly once. Because the cost of generating each such arc, and the cost of checking if this is in conflict with a chosen tuple are $O(1)$, the total cost is $O(e)$.

Complexity of single-source reachability algorithm: Example 4.4

This case is very similar to the previous one. The size of `reach` is bounded by $O(n)$, and so is the size of the `chosen` and `theta` relations. However, in the process of generating `reach`, all the edges reachable from the source node are explored by the algorithm exactly once. Thus the worst case complexity is $O(e)$.

Complexity of simple path: Example 4.5

Again, all the edges in the graph will be visited in the worst case, yielding complexity $O(e)$, where $e = n^2$, according to our assumption that the graph is complete. Every arc is visited once and, therefore, the global complexity is $O(e)$, with $e = n^2$.

Complexity of a bipartite matching: Example 4.2

Initially all body tuples are inserted into the `theta` relation at cost $O(e)$. The computation terminates in $O(\min(n_1, n_2)) = O(n)$ steps, where n_1 and n_2 are, respectively, the number of nodes in the left and right parts of the graph. At each step, one tuple t is selected at cost $O(1)$ and the tuples conflicting with the selected tuple are deleted. The global cost of deleting conflicting tuples is $O(e)$, since we assume that each tuple is accessed in constant time. Therefore the global cost is $O(e)$.

Complexity of linear sequencing of the elements of a set: Example 4.1

If n is the cardinality of the domain d , the computation terminates in $O(n)$ steps. At each step, n tuples are computed, one tuple is chosen and the remaining tuples are discarded. Therefore the complexity is $O(n^2)$.

7.3 Implementation of choice-least/most programs

In the presence of choice-least (or choice-most) goals, the best alternative must be computed, rather than an arbitrary one chosen at random. Thus, we introduce a new algorithm, reported in Figure 2, derived from the algorithm of Figure 1.

Let us consider the general case where programs could contain three different kinds of choice rules: (i) choice-least rules that have one choice least goal, and zero or more choice goals; (ii) choice-most rules that have one choice-most goal and zero or more choice goals; and (iii) pure choice rules that have one or more choice goals and no choice-least or choice-most goals. Then Step 3 (i) in Algorithm 7.1 should be modified as follows:

- 3: (i) If r is a choice-least (choice-most) rule then select a single tuple $x \in \text{theta}_r$ with least (most) cost; otherwise (r is pure choice rule, so) take an arbitrary $x \in \text{theta}_r$. Move x from `thetar` to `chosenr`;

An additional optimization is however possible, as discussed next. Consider for instance Prim's algorithm in Example 6.2:

Say that `theta` contains two tuples $t_1 = (x, y_1, c_1)$ and $t_2 = (x, y_2, c_2)$. Then, the following properties hold for Algorithms 1 with Step 3 (i) modified as shown above:

- If $c_1 < c_2$ then t_2 is not a least-cost tuple,
- t_1 belongs to $\text{conflict}(\text{theta}, \delta)$, if and only if t_2 does.

Therefore, the presence of t_2 is immaterial to the result of the computation, and we can modify our algorithm to ensure that only t_1 is kept in table theta . This improvement can be implemented by ensuring that the attribute Y is unique key for the table theta . When a new tuple t' is generated and a tuple with the same key value is found in theta , the tuple with the smaller cost value is entered in the table and the other is discarded. This reduces, the maximum cardinality of theta for Prim's and Dijkstra's algorithm from e (number of arcs) to n , (number of nodes).

However, the above considerations are not valid for rules containing more than one choice atoms. For instance, in the greedy TSP program, or the optimal matching program (Examples 6.5 and 6.1, respectively), the choice rules have the following choice goals:

$\text{choice}((X), (Y)), \text{choice}((Y), (X)), \text{choice-least}((Y), (C))$

Say that theta contains the following tuples: $t_1 = (x_1, y_1, c_1)$, $t_2 = (x_1, y_2, c_2)$ $t_3 = (x_2, y_2, c_3)$, with $c_1 < c_2 < c_3$. Although, c_3 conflicts with c_2 and has larger cost value, it cannot be eliminated, since it has a chance to be selected later. For instance, if t_1 is selected first then t_2 will be eliminated, since it conflicts with t_1 . But t_3 does not conflict with t_1 , and remains, to be selected next.

Thus, the general rule is as follows:

1. The union of the left sides of all choice goals is a unique key for theta .
2. When a new tuple is inserted and there is a conflict on the unique key value, retain in theta only the tuple with the lesser cost.

The above optimization can be carried out by modifying Step 1 (iii) in Algorithm 7.1 as follows:

- 1: (iii) Add each tuple of θ_r to theta_x ; when key conflicts occur, and r is a choice-least (choice-most) table, retain the lesser (larger) of the tuples.

Moreover, insertion and deletion of an element from a theta table can be done in constant time, since we are assuming that hash indexes are available, whereas the selection of a least/most cost element is done in linear time. The selection of the least/most cost element can be done in constant time by organizing theta as priority queues. However the cost of insertion, deletion of least/most cost element from a priority queue are now logarithmic, rather than constant time. Therefore, when using priority queues we can improve the performance of our algorithm by delaying the merging θ_r into theta_x (Step 2), as to allow the elimination from θ_r of tuples conflicting with the new δ selected at Step 3. Therefore, we will move one tuple from θ_r to theta_x (if $\theta_r \neq \emptyset$), in Step 1 (iii), and the remaining tuples at the end of Step 3 (those conflicting with δ excluded).

Observe that the improvement performed in Step 3 reduces the max cardinality of tables theta_x . In our Prim's algorithm the size of the table theta is reduced from the number of edges e to the number of nodes n . This has a direct bearing on the performance of our algorithm since the selection of a least cost tuple is performed

Algorithm 7.2

Greedy Semi-naive computation of a choice model.

Input: Choice program P .**Output:** I , a greedy choice model for P .**begin**

0: Initialization.

For every $r \in P'_C$ set $\text{chosen}_r = \text{theta}_r = \emptyset$;Set: $I := T_{P'_D}^{\uparrow\omega}(\emptyset)$;1: **Repeat**(i) Select an unmarked arbitrary $r \in P'_C$ and mark r (ii) Compute: $\theta_r = (T_r(I) \setminus \text{theta}_r) \setminus \text{conflict}(T_r(I) \setminus \text{theta}_r, \text{chosen}_r)$;(iii) If r is a choice-least (choice-most) rule then select a single tuple $x \in \theta_r$ with least (most) cost; otherwise (r is pure choice rule, so) take an arbitrary $x \in \theta_r$. Move x from θ_r to theta_r ;**Until** $\text{theta}_r \neq \emptyset$ or all rules in P'_C are marked;2: **If** $\text{theta}_r = \emptyset$ **Return** I ;3: (i) If r is a choice-least (choice-most) rule then select a single tuple $x \in \text{theta}_r$ with least (most) cost; otherwise (r is pure choice rule and) take an arbitrary $x \in \text{theta}_r$. Move x from theta_r to chosen_r ;(ii) With $\delta = \{x\}$, delete from the selected table theta_r every tuple in $\text{conflict}(\text{theta}_r, \delta)$.(iii) Add each tuple of θ_r to theta_r ; when key conflicts occur, and r is a choice-least (choice-most) table, retain the lesser (larger) of the tuples.4: Set: $I = T_{P'_D}^{\uparrow\omega}(I \cup \delta)$, using the differential fixpoint improvement, then unmark all P'_C rules and resume from Step 1.**end.**

Fig. 2. Greedy semi-naive computation.

n times. Now, if a linear search is used to find the least-cost element the global complexity is $O(n \times n)$. Similar considerations and complexity measures hold for Dijkstra's algorithm.

In some cases, however, the unique key improvement just describe might be of little or no benefit. For the TSP program and the optimal matching program, where the combination of both end-points is the key for theta , no benefit is to be gained since there is at most one edge between the two nodes. (In a database environment this might follow from the declaration of unique keys in the schema, and can thus be automatically detected by a compiler.)

Next, we compute the complexity of the various algorithms, assuming that the theta tables are supported by simple hash-based indexes, but there is no priority queue. The complexities obtained with priority queues are discussed in the next section.

Complexity of Prim's Algorithm: Example 6.2

The computation terminates in $O(n)$ steps. At each step, $O(n)$ tuples are inserted into the table theta , one least-cost tuple is moved to the table chosen and conflicting

tuples are deleted from θ . Insertion and deletion of a tuple is done in constant time, whereas selection of the least cost tuple is done in linear time. Since the size of θ is bounded by $O(n)$, the global complexity is $O(n^2)$.

Complexity of Dijkstra algorithm: Example 6.3

The overall cost is $O(n^2)$ as for Prim's algorithm.

Complexity of sorting the elements of a relation: Example 6.4

At each step, n candidate tuples are generated, one is chosen, and all tuples are eliminated from θ . Here, each new X from succ is matched with every Y , even when differential techniques are used. Therefore, the cost is $O(n^2)$.

Greedy TSP: Example 6.5

The number of steps is equal to n . At each step, n tuples are computed by the evaluation of the body of the chosen rule and stored into the temporary relation θ . Then, one tuple with least cost is selected (Step 1) and entered in θ all the remaining tuples are deleted from the relation (Step 3). The cost of inserting one tuple into the temporary relation is $O(1)$. Therefore, the global cost is $O(n^2)$.

Optimal matching in a directed graph: Example 6.1

Initially all body tuples are inserted into the θ relation at cost $O(e)$. The computation terminates in $O(n)$ steps. At each step, one tuple t with least cost is selected at cost $O(e)$ and the tuples conflicting with the selected tuple are deleted. The global cost of deleting conflicting tuples is $O(e)$ (they are accessed in constant time). Therefore the global cost is $O(e \times n)$.

7.4 Priority queues

In many of the previous algorithms, the dominant cost is finding the least value in the table $\theta_{\mathbf{r}}$, where \mathbf{r} is a least-choice or most-choice rule. Priority queues can be used to reduce the overall cost.

A priority queue is a partially ordered table where the cost of the i th element is greater or equal than the cost of the $(i \div 2)$ th element (Aho *et al.*, 1974). Therefore, our table $\theta_{\mathbf{r}}$ can be implemented as a list where each node having position i in the list also contains (1) a pointer to the next element, (2) a pointer to the element with position $2 \times i$, and (3) a pointer to the element with position $i \div 2$. The cost of finding the least value is constant-time in a priority queue, the cost of adding or deleting an element is $\log(m)$ where m is the number of the entries in the queue.

Also, in the implementation of Step 2 (ii), a linear search can be avoided by adding *one search index for each left side of a choice or choice-least goal*. For instance, for Dijkstra's algorithm there should be a search index on X , for Prim's on Y . The

operation of finding the least cost element in θ_r can be done during the generation of the tuples at no additional cost. Then we obtain the following complexities:

Complexity of Prim's algorithm: Example 6.2

The computation terminates in $O(n)$ steps and the size of the priority queue is bounded by $O(n)$. The number of candidate tuples is bounded by $O(e)$. Therefore, the global cost is bounded by $O(e \times \log n)$.

Complexity of Dijkstra algorithm: Example 6.3

The overall cost is $O(e \times \log n)$ as for Prim's algorithm.

Complexity of sorting the elements of a relation: Example 6.4

The number of steps is equal to n . At each step, n tuples are computed, one is stored into θ and next moved to θ_r while all remaining tuples are deleted from θ . The cost of each step is $O(n)$ since deletion of a tuple from θ is constant time. Therefore, the global cost is $O(n^2)$.

Greedy TSP: Example 6.5

Observe that θ_r here contains at most one tuple. The addition of the first tuple into an empty priority queue, θ_r , and the deletion of the last tuple from it are constant time operations. Thus the overall cost is the same as that without a priority queue, i.e. the global cost is $O(n^2)$.

Optimal matching in a bipartite graph: Example 6.1

Initially, all body tuples are inserted into the θ relation at cost $O(e \times \log e)$. The computation terminates in $O(n)$ steps. At each step, one tuple t is selected and all remaining tuples conflicting with t are deleted (the conflicting tuples here are those arcs having the same node as source or end node of the arc). The global number of extractions from the priority queue is $O(e)$. Therefore, the global complexity is $O(e \times \log e)$.

Observe that, using a priority queues, an asymptotically optimum performance (Aho *et al.*, 1974) has been achieved for all problems, but that of sorting the elements of a domain, Example 6.4. This problem is considered in the next section.

7.5 Discussion

A look at the structure of the program in Example 6.4 reveals that at the beginning of each step a new set of (x, y) pairs is generated for θ by the two goals

$\text{succ}(_, X), d(Y)$ which define a Cartesian product. Thus, the computation can be represented as follows:

$$\Theta = \pi_2 \text{succ} \times d$$

where succ and d are the relations containing their homonymous predicates. We can also represent θ and theta as Cartesian products:

$$\theta = (\pi_2 \delta \times d) \setminus \text{chosen} = \pi_2 \delta \times (d \setminus \pi_2 \text{chosen})$$

Therefore, the key to obtaining an efficient implementation here consists in storing only the second column of the theta relation, i.e.

$$\pi_2 \text{theta} = d \setminus \pi_2 \text{chosen}$$

The operation of selecting a least-cost tuple from theta now reduces to that of selecting a least-cost tuple from $\pi_2 \text{theta}$, which therefore should be implemented as a priority queue.

Assuming these modifications, we can now recompute the complexity of our Example 3, by observing that all the elements in d are added to $\pi_2 \text{theta}$ once at the first iteration. Then each successive iteration eliminates one element from this set. Thus, the overall complexity is linear in the number of nodes. For Example 11, the complexity is $O(n \times \log n)$ if we assume that a priority queue is kept for $\pi_2 \text{theta}$. Thus we obtain the optimal complexities.

No similar improvement is applicable to the other examples, where the rules do not compute the Cartesian product of two relations. Thus, this additional improvement could also be incorporated into a smart compiler, since it is possible to detect from the rules whether theta is in fact the Cartesian product of its two subprojections. However this is not the only alternative since many existing deductive database systems provide the user with enough control to implement this, and other differential improvements previously discussed, by coding them into the program. For instance, the $\mathcal{L}\mathcal{D}\mathcal{L}++$ users could use XY-stratified programs for this purpose (Zaniolo *et al.*, 1993); similar programs can be used in other systems (Vaghani *et al.*, 1994).

8 Conclusion

This paper has introduced a logic-based approach for the design and implementation of greedy algorithms. In a nutshell, our design approach is as follows: (i) formulate the all-answer solution for the problem at hand (e.g. find all the costs of all paths from a source node to other nodes); (ii) use choice-induced FD constraints to restrict the original logic program to the nondeterministic generation of a single answers (e.g. find a cost from the source node to each other node); and (iii) specialize the choice goals with preference annotations to force a greedy heuristics upon the generation of single answers in the choice-fixpoint algorithm (thus computing the least-cost paths). This approach yields conceptual simplicity and simple programs; in fact it has been observed that our programs are often similar to pseudo code expressing the same problem in a procedural language. However, our approach offers additional

advantages, including a formal logic-based semantics and a clear design method, implementable by a compiler, to achieve optimal implementations for our greedy programs. This method is based on

- The use of chosen tables and theta tables, and of differential techniques to support the second kind of table as a concrete view. The actual structure of theta tables, their search keys and unique keys are determined by the choice and choice-least goals, and the join dependencies implied by the structure of the original rule.
- The use of priority queues for expediting the finding of extrema values.

Once these general guidelines are followed (by a user or a compiler) we obtain an implementation that achieves the same asymptotic complexity as procedural languages.

This paper provides a refined example of the power of Kowalski's seminal idea: *algorithms = logic + control*. Indeed, the logic-based approach here proposed covers all aspects of greedy algorithms, including (i) their initial derivation using rules with choice goals, (ii) their final formulation by choice-least/most goals, (iii) their declarative stable-model semantics, (iv) their operational (fixpoint) semantics, and finally, (v) their optimal implementation by syntactically derived data structures and indexing methods. This vertically integrated, logic-based, analysis and design methodology represents a significant step forward with respect to previous logic-based approaches to greedy algorithms (including those we have proposed in the past (Greco *et al.*, 1992; Ganguly *et al.*, 1995)).

Acknowledgements

The authors would like to thank D. Saccà for many helpful discussions and suggestions. The referees deserve credit for many improvements.

References

- Abiteboul, S., Hull, R. & Vianu, V. (1994) *Foundations of Databases*. Addison-Wesley.
- Abiteboul, S. & Vianu, V. (1991) Datalog Extensions for Databases Queries and Updates. *J. Comput. & Syst. Sci.*, **43**(1), 62–124.
- Aho, A. V., Hopcroft, J. E. & Ullman, J. D. (1974) *The Design and Analysis of Computer Algorithms*. Addison-Wesley.
- Apt, K. R., Blair, H. A. & Walker, A. (1988) Towards a theory of declarative knowledge, In: J. Minker (editor), *Foundations of Deductive Databases and Logic Programming*, pp. 89–148. Morgan Kaufmann.
- Chandra, A. & Harel, D. (1982) Structure and Complexity of Relational Queries, *J. Comput. & Syst. Sci.*, **25**(1), 99–128. Academic Press.
- Dietrich, S. W. (1992) Shortest Path by Approximation in Logic Programs, *ACM Lett. Programming Lang. & Syst.*, **1**(2), 119–137.
- Ganguly, S., Greco, S. & Zaniolo, C. (1995) Extrema Predicates in deductive databases. *J. Comput. & Syst. Sci.*, **51**(2), 244–259.
- Gelfond, M. & Lifschitz, V. (1988) The stable model semantics of logic programming. *Proc. 5th Int. Conf. on Logic Programming*, pp. 1070–1080. Jerusalem, Israel.

- Giannotti, F., Pedreschi, D., Saccà, D. & Zaniolo, C. (1991) Nondeterminism in deductive databases. *Proc. 2nd Int. Conf. on Deductive and Object-Oriented Databases*, Munich, Germany. (Lecture Notes in Computer Science 566, pp. 129–146. Springer-Verlag.)
- Giannotti, F., Pedreschi, D. & Zaniolo, C. (2000) Semantics and expressive power of non-deterministic constructs in deductive databases. *J. Comput. & Syst. Sci.* (to appear).
- Greco, S. & Zaniolo, C. (1999) Dynamic programming in Datalog with aggregates. *IEEE Trans. Knowl. Eng.*, **11**(2), 265–283.
- Greco, S., Saccà, D. & Zaniolo, C. (1995) DATALOG queries with stratified negation and choice: from P to D^P . *Proc. 5th Int. Conf. on Database Theory: Lecture Notes in Computer Science 893*, pp. 81–96. Springer-Verlag.
- Greco, S., Zaniolo, C. & Ganguly, S. (1992) Greedy by Choice. *Proc. 11th ACM Symposium on Principles of Database Systems*, pp. 105–113. San Diego, CA. ACM Press.
- Greco, S. & Zaniolo, C. (1998) Greedy algorithms in Datalog with choice and negation. *Proc. Int. Joint Conf. and Symposium on Logic Programming*, pp. 294–309. Manchester, UK. MIT Press.
- Krishnamurthy, R. & Naqvi, S. (1988) Non-deterministic choice in Datalog. *Proc. 3rd Int. Conf. on Data and Knowledge Bases*, pp. 416–424. Jerusalem, Israel. Morgan Kaufmann.
- Marek, W. & Truszczynski, M. (1991) Autoepistemic logic. *J. ACM*, **38**(3), 588–619.
- Moret, B. M. E. & Shapiro, H. D. (1993) *Algorithms from P to NP*. Benjamin Cummings.
- Papadimitriou, C. & Steiglitz, K. (1975) *Combinatorial Optimization: Algorithms and Complexity*. Prentice Hall.
- Przymusiński, T. (1988) On the declarative and procedural semantics of stratified deductive databases. In: J. Minker (editor), *Foundations of Deductive Databases and Logic Programming*, pp. 193–216. Morgan-Kaufman.
- Przymusińska, A. & Przymusiński, T. (1988) Weakly perfect model semantics for logic programs. *Proc. 5th Int. Conf. and Symposium on Logic Programming*, pp. 1106–1122. Seattle, WA. MIT Press.
- Ross, K. A. & Sagiv, Y. (1992) Monotonic aggregation in deductive databases. *Proc. 11th ACM Symposium on Principles of Database Systems*, pp. 114–126. San Diego, CA. ACM Press.
- Saccà, D. & Zaniolo, C. (1990) Stable models and non-determinism in logic programs with negation. *Proc. 9th ACM Symposium on Principles of Database Systems*, pp. 205–217. Nashville, TN. ACM Press.
- Sudarshan, S. & Ramakrishnan, R. (1991) Aggregation and relevance in deductive databases. *Proc. 17th Conf. on Very Large Data Bases*, pp. 501–511. Barcelona, Spain. Morgan Kaufmann.
- Ullman, J. D. (1989) *Principles of Data and Knowledge-Base Systems*, Vol. 1 & 2. Computer Science Press.
- Vaghani, J., Ramamohanarao, K., Kemp, D. B., Somogyi, Z., Stuckey, P. J., Leask, T. S. & Harland, J. (1994) The Aditi deductive database system. *The VLDB J.*, **3**(2), 245–288.
- Van Gelder, A., Ross, K. A. & Schlipf, J. S. (1991) The well-founded semantics for general logic programs. *J. ACM*, **38**(3), 620–650.
- Van Gelder, A. (1993) Foundations of aggregations in deductive databases. *Proc. 3rd Int. Conf. on Deductive and Object-Oriented Databases*, Phoenix, AZ. (Lecture Notes in Computer Science 760, pp. 13–34. Springer-Verlag.)
- Zaniolo, C., Arni, N. & Ong, K. (1993) Negation and aggregates in recursive rules: the $\mathcal{L}\mathcal{Q}\mathcal{L}++$ approach. *Proc. 3rd Int. Conf. on Deductive and Object-Oriented Databases*, Phoenix, AZ. (Lecture Notes in Computer Science 760, pp. 204–221. Springer-Verlag.)
- Zaniolo C., Ceri, S., Faloutsos, C., Subrahmanian, V. S. & Zicari, R. (1997) *Advanced Database Systems*. Morgan Kaufmann.