

Logic Programming and Deductive Databases

Chapter 13: Magic Sets

Prof. Dr. Stefan Brass

Martin-Luther-Universität Halle-Wittenberg

Summer 2022

<http://www.informatik.uni-halle.de/~brass/lp22/>

Objectives

After completing this chapter, you should be able to:

- explain the basic idea of the magic set transformation.
- What does a magic fact mean?
- perform the magic set transformation for a given input program.
- name and explain different SIP strategies.

SIP: "Sideways Information Passing".

Contents

- ① Introduction
- ② SIP-Strategies, Adorned Program
- ③ The Magic Set Transformation
- ④ Improvements

Introduction (1)

Example (Grandparents of Julia):

- Logic Program (IDB-Predicates and Query):

```
parent(X, Y)      ← mother(X, Y).  
parent(X, Y)      ← father(X, Y).  
grandparent(X, Z) ← parent(X, Y) ∧  
                    parent(Y, Z).  
answer(X)         ← grandparent(julia, X).
```

- EDB-Predicates (stored in the database):

- father
- mother

Introduction (2)

Problem:

- Naive/Seminaive Bottom-Up Evaluation computes all parent- and grandparent-relationships of all persons in the database.
- Until now, the actual query is considered only at the very end of query evaluation — after the entire minimal model was computed.
- Therefore, the method is not goal-directed: It computes many superfluous facts, which are not relevant for the query.

Introduction (3)

Solution:

- The “Magic Set” Transformation rewrites the program such that the rules can only “fire” when their result (the derived fact) is relevant for the query.
- This is done by making the occurring queries and subqueries explicit. They are encoded as facts of “magic predicates”.
- E.g. the query

`? grandparent(julia,X)`

is represented as

`m_grandparent_bf(julia).`

Introduction (4)

More About Encoding of Queries:

- Consider again the correspondance:
 - Query: `? grandparent(julia,X)`
 - Magic Fact: `m_grandparent_bf(julia).`
- Magic facts should be representable in a database, and therefore should not contain variables.
- Solution: The binding pattern indicates the position of the variables (their name is not important).
- Only the values of constants in the query (bound arguments) are explicitly stored in the magic fact.

Introduction (5)

Example Output, First Part:

- Rules are restricted by an additional body literal so that they can fire only if there is a matching query:

```
parent(X, Y)      ←  m_parent_bf(X) ∧  
                    mother(X, Y).  
parent(X, Y)      ←  m_parent_bf(X) ∧  
                    father(X, Y).  
grandparent(X, Z) ←  m_grandparent_bf(X) ∧  
                    parent(X, Y) ∧  
                    parent(Y, Z).  
answer(X)         ←  true ∧  
                    grandparent(julia, X).
```


Introduction (6)

Example Output, Second Part:

```
m_grandparent_bf(julia) ← true.  
m_parent_bf(X)           ← m_grandparent_bf(X).  
m_parent_bf(Y)           ← m_grandparent_bf(X) ∧  
                           parent(X, Y).
```

- Of course, the original query must be represented as a magic fact in the rewritten program.
- In addition, magic facts corresponding to the occurring subqueries must be derivable.
- Example: To compute the grandparents of X , one must first compute the parents of X .

Introduction (7)

Summary (Equivalence):

- The “Magic Set” transformation produces a program which is equivalent for the given query:
 - The extension of the predicate “`answer`” in the minimal model of the transformed program is the same as in the minimal model of the original program.
- But often the minimal model of the transformed program is much smaller than the minimal model of the original program.

It contains only IDB-facts that are relevant for the given query.

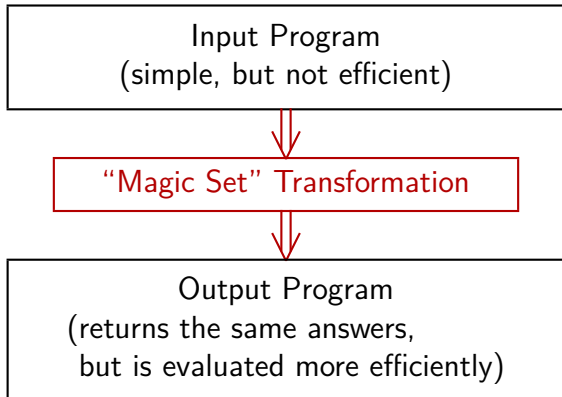
Introduction (8)

Summary (Equivalence), Continued:

- This equivalence (for the predicate `answer`) is independent of the extensions of the EDB-predicates:
 - No database access is needed during the transformation.

Therefore, the transformation itself is quite efficient (one usually assumes that external memory accesses are expensive).
 - The transformed program can be executed several times, even when the database state was changed in the meantime.

Introduction (9)



Introduction (10)

Magic Sets and Built-in Predicates:

- The magic set transformation can also improve the termination of programs with built-in predicates:
 - E.g., `append` has an infinite extension.
 - But for a concrete query, only a finite number of facts might be needed.

E.g., when two given lists are appended.

- In this way, the magic set transformation might turn an infinite minimal model into a finite one.

Of course, it depends on the program and the query whether this works.

Introduction (11)

Top-Down vs. Bottom-Up:

- Before the magic set transformation, there were two competing evaluation approaches:
 - Top-down evaluation (e.g. SLD-resolution):
Starts from the query, simplifies it, until facts can be used. Advantage: Goal-directed.
 - Bottom-up evaluation (e.g., seminaive method):
Starts from the facts, computes derived facts, until answers to the query are reached. Advantage: Avoids duplicate work, ensures termination.
- Magic sets combine the advantages of both.

Introduction (12)

Magic Sets — A Source Code Level Transformation:

- The underlying “bottom-up machine” can remain unchanged.
- It is always good to separate problems and solve them one after the other.
- One can understand the method on a high level of abstraction (Herbrand models instead of internal data structures).
- However, it is probably advantageous for an implementation to treat the magic predicates specially.

Introduction (13)

IDB predicates as procedures:


- With magic sets, IDB predicates can be understood as procedures with input and output arguments:
 - Input: Relation M with bindings for the input arguments (this is the “Magic Set”).
 - Output: Relation R for all arguments.

It does not suffice to return only a relation for the output arguments because the connection to the input arguments would not be clear if M contains more than one tuple.
- If E is the original extension of the predicate, $R = E \bowtie M$ holds (this is a semi-join).

Introduction (14)

IDB predicates as procedures, continued:

- Example:



The database might contain many more mother/father relationships, but only the required parent tuples are derived.

- However, for recursive predicates, *M* might still be extended later.

Contents

- 1 Introduction
- 2 SIP-Strategies, Adorned Program
- 3 The Magic Set Transformation
- 4 Improvements

SIP-Strategies (1)

Motivation:

- SIP = „Sideways Information Passing“

When rules are evaluated, information is passed “sideways”: from a body literal that is evaluated earlier to one that is evaluated later.

- $\text{grandparent}(X, Z) \leftarrow \text{parent}(X, Y) \wedge \text{parent}(Y, Z).$
 - Variable binding from the caller: $X = \text{julia}.$
 - This is passed to the first body literal.
 - By evaluating this body literal, one gets bindings for Y , e.g. $Y = \text{eric}$ and $Y = \text{fiona}.$
 - These are passed to the second body literal.

SIP-Strategies (2)

- However, for the query “? **grandparent**(X, **alan**)”, it is more efficient to start the evaluation of

$$\text{grandparent}(X, Z) \leftarrow \text{parent}(X, Y) \wedge \text{parent}(Y, Z)$$

with the second body literal.

- Then the binding $Z = \text{alan}$ can be used.
 - This gives bindings for Y , which can be passed to the first body literal.
- If instead one evaluates the first body literal first, this is done with the binding pattern `ff`, and one has to compute the complete extension of `parent`.

SIP-Strategies (3)

Definition:

- Given a rule $A \leftarrow B_1 \wedge \dots \wedge B_m$ and a binding pattern β for $\text{pred}(A)$,
- a SIP-strategy defines an evaluation sequence for the body literals, i.e. a permutation

$$\pi: \{1, \dots, m\} \rightarrow \{1, \dots, m\},$$

- and for every $k \in \{1, \dots, m\}$ a valid binding pattern $\beta_{\pi(k)} \in \text{valid}(\text{pred}(B_{\pi(k)}))$ such that

$$\text{input}(B_{\pi(k)}, \beta_{\pi(k)}) \subseteq \text{input}(A, \beta) \cup \bigcup_{j=1}^{k-1} \text{vars}(B_{\pi(j)}).$$

SIP-Strategies (4)

Note:

- This is the same condition for π and $\beta_{\pi(k)}$ as in the definition of “range restricted rule”.
- A given evaluation sequence determines “maximally bound” binding patterns for the body literals:
 - Values for variables in “bound” argument positions in the head literal are known.
 - Values for variables in body literals that were evaluated earlier are known.
 - All other variables do not have a known value yet, thus they lead to “free” argument positions.

SIP-Strategies (5)

- Most SIP-strategies choose the above “maximal” binding pattern that uses all existing bindings.
- Therefore, the real decision is the evaluation sequence for the body literals. The binding patterns are then often automatically determined.

However, the possible evaluation sequences depend on the valid binding patterns for the body literal: Some predicates can only be evaluated if certain arguments are bound.

- A SIP-strategy can ignore existing bindings and choose a more general binding pattern.

Possible reasons are explained on the next slide.

SIP-Strategies (6)

Reasons for not choosing the maximal binding pattern:

- Not all binding patterns might be implemented.

This is obvious for built-in predicates, but happens also for IDB-predicates in separately compiled modules (modules define “exported binding patterns”).

- One needs the more general binding pattern anyway at some other place in the program.

In this way, one avoids duplicating the rules. But unless the other binding pattern is $ff \dots f$, one computes more tuples (if calls disjoint).

- to simplify the magic rules.

An even more general kind of “SIP-strategy” permits to choose a subset of the earlier evaluated body literals for the magic rule.

SIP-Strategies (7)

Exercise:

- Consider this rule (with only IDB predicates):

```
ship_to(ProdName, City) ←  
    has_ordered(CustNo, ProdNo) ∧  
    customer_city(CustNo, City) ∧  
    product_name(ProdNo, ProdName).
```

- Select a good evaluation sequence for each of the following calls, and state the binding patterns:
 - `ship_to(X, 'Halle').`
 - `ship_to('Van Tastic', X).`
 - `ship_to('Van Tastic', 'Halle').`

SIP-Strategies (8)

Framework for SIP-Strategies:

- Let $A \leftarrow B_1 \wedge \dots \wedge B_m$ be called with binding pattern β . One chooses first $\pi(1)$ (i.e. the first body literal to evaluate), then $\pi(2)$, and so on.
- Given that one has already chosen $\pi(1), \dots, \pi(k)$, the i -th body literal is possible with binding pattern $\beta' \in \text{valid}(\text{pred}(B_i))$ iff the literal has not been chosen yet, i.e. $i \in \{1, \dots, m\} \setminus \{\pi(1), \dots, \pi(k)\}$ and

$$\text{input}(B_i, \beta') \subseteq \text{input}(A, \beta) \cup \bigcup_{j=1}^k \text{vars}(B_{\pi(j)}).$$

SIP-Strategies (9)

Example:

- Consider the following rule is called with $p(X, 3)$:

$$p(X, Y) \leftarrow X < Y \wedge q(X).$$

- At the beginning, only the second body literal is evaluable (with binding pattern \mathbf{f}).

The only valid binding pattern for $<$ is \mathbf{bb} , therefore the first body literal cannot be evaluated at this point (although it has more bound arguments than $q(X)$: the value for Y is already known).

- Thus, all SIP-strategies must select $\pi(1) = 2$.
- After $q(X)$ is evaluated, the value of X is known, and the first literal becomes evaluable: $\pi(2) = 1$.

SIP-Strategies (10)

Common SIP-Strategies:

- Among all possible (i, β) , choose one such that β has the smallest number of free argument positions.
- Among all possible (i, β) , choose one such that β has the largest number of bound arguments.
- Among all possible (i, β) , choose one such that i is minimal.

This strategy evaluates body literals in the sequence given by the programmer as far as possible. If there should be several possible β for the minimal i , choose one with the maximum number of bound argument positions.

SIP-Strategies (11)

Example:

- Consider the rule

$$p(X_1, X_2) \leftarrow q(X_1, Y) \wedge r(X_1, X_2, Z_1, Z_2)$$

and the call $p(a, b)$.

- A SIP-strategy that tries to maximize the number of bound argument positions begins the evaluation with the second body literal: $r(X_1, X_2, Z_1, Z_2)$.
- A SIP-strategy that minimizes the number of free argument positions chooses $q(X_1, Y)$ first.

SIP-Strategies (12)

Further Input for SIP-Strategies:

- The SIP-strategy is an important component of the optimizer of a deductive DBMS.
- It should take also the following information into account:
 - Keys
 - Indexes
 - Size of Relations
 - Number of different values in an attribute

SIP-Strategies (13)

Goals of SIP-Strategies:

- Try to keep intermediate results small.
- “Fail as early as possible”.
- Call expensive predicates only when cheap tests were successful. Expensive predicates are:
 - Recursive predicates
 - Predicates that need themselves many joins and possibly a duplicate elimination.
 - Built-in predicates with complicated computations or slow network accesses.

SIP-Strategies (14)

Exercises:

- $p(X, Z) \leftarrow q(X, Y_1, Y_2, Y_3) \wedge r(Y_1, Z)$, with call $p(a, b)$, when the first argument of q and r is a key.
- Suppose p is defined by $p(X) \leftarrow q_1(X) \wedge q_2(X)$ and is called with binding pattern \mathbf{f} . Cost estimates:
 - $q_1: \mathbf{f}$ produces 100 tuples in 100 ms.
 - $q_1: \mathbf{b}$ checks a single value in 3 ms (index).
 - $q_2: \mathbf{f}$ produces 1000 tuples in 200 ms.
 - $q_2: \mathbf{b}$ checks a single value in 100 ms (FT scan).
 - Sorting/intersecting the two sets costs 1000 ms.

Contents

- 1 Introduction
- 2 SIP-Strategies, Adorned Program
- 3 The Magic Set Transformation**
- 4 Improvements

Adorned Program (1)

Goal:

- The first step of the transformation is to make the binding patterns of the IDB-predicates explicit.

This simplifies the definition of the main magic set transformation.

- I.e. the predicate p is renamed to p_β . Several versions of a predicate are produced for different β .

Sometimes one IDB-predicate is called with different binding patterns.

- Furthermore, the body literals are ordered in the evaluation sequence.

In this way, the information from the SIP-strategy is encoded in the program.

Adorned Program (2)

Definition:

- Let a logic program P , valid binding patterns $valid$ and a SIP-strategy \mathcal{S} be given.
- For each predicate $p \in \mathcal{P}_{IDB}(P)$, $p \neq \text{answer}$, and every binding pattern $\beta \in valid(p)$ a new predicate $p_{-}\beta$ is introduced.
- For a literal $A = p(t_1, \dots, t_n)$ and binding pattern $\beta \in valid(p)$ let

$$adorn(A, \beta) := \begin{cases} p_{-}\beta(t_1, \dots, t_n) & \text{if } p \in \mathcal{P}_{IDB}(P), \\ & p \neq \text{answer} \\ p(t_1, \dots, t_n) & \text{otherwise.} \end{cases}$$

Adorned Program (3)

Definition, continued:

- The program $AD(P)$ contains for each rule

$$A \leftarrow B_1 \wedge \cdots \wedge B_m \quad \text{from } IDB(P)$$

and each binding pattern $\beta \in \text{valid}(\text{pred}(A))$
the rule

$$\text{adorn}(A, \beta) \leftarrow \text{adorn}(B_{\pi(1)}, \beta_{\pi(1)}) \wedge \cdots \wedge \text{adorn}(B_{\pi(m)}, \beta_{\pi(m)}),$$

where π and β_1, \dots, β_m are determined by the SIP-strategy \mathcal{S} .

Adorned Program (4)

Note:

- It is theoretically simpler to process every rule for each possible binding pattern.
- The unnecessary binding patterns are eliminated when one later deletes all predicates (and their definitions) on which “**answer**” does not depend.
- In practice, only necessary $p_{-\beta}$ are constructed.

E.g. one manages a set of all combinations (p, β) that still have to be processed. This is initialized with $\{(\text{answer}, f \dots f)\}$. In each step, one takes an element from this set and generates the rules for $p_{-\beta}$. If the rule bodies contain new combinations of predicates and binding patterns, one inserts them into the set.

Adorned Program (5)

- Binding patterns chosen by the SIP-strategy for EDB-/builtin predicates are important to determine which index/implementation variant is to be used.
- But the names of these predicates are determined in the database/the system, they cannot be changed.

In contrast, names of IDB-predicates (\neq `answer`) are only important within the program.

- Furthermore, the explicit binding patterns are a preparation for the magic set transformation. This is not useful for EDB-predicates, because their complete extensions are already stored.

Adorned Program (6)

Exercise:

- Compute $AD(P)$ (as far as relevant for **answer**).
Use a good SIP-strategy:
 - $\text{parent}(X, Y) \leftarrow \text{mother}(X, Y).$
 $\text{parent}(X, Y) \leftarrow \text{father}(X, Y).$
 $\text{grandparent}(X, Z) \leftarrow \text{parent}(X, Y) \wedge \text{parent}(Y, Z).$
 $\text{answer}(X) \leftarrow \text{grandparent}(\text{julia}, X).$
 - $\text{append}([], L, L) \leftarrow \text{true}.$
 $\text{append}(X, L, Y) \leftarrow \text{cons}(H, T, X) \wedge \text{cons}(H, TL, Y) \wedge \text{append}(T, L, TL).$
 $\text{answer}(X) \leftarrow \text{append}([a], [b, c], X).$

Adorned Program (7)

Exercise, continued:

- Compute $AD(P)$ (as before):
 - „Same Generation Cousins“:
 $sg(X, X) \leftarrow person(X).$
 $sg(X, Y) \leftarrow parent(X, Xp) \wedge parent(Y, Yp) \wedge sg(Xp, Yp).$
 $answer(X) \leftarrow sg(julia, X).$
 - Abstract Example:
 $answer(yes) \leftarrow p(a, b).$
 $p(X, Z) \leftarrow q(X, Y) \wedge q(Y, Z).$
 $q(X, Y) \leftarrow r(a, X, Y).$

Adorned Program (8)

Lemma:

- Let \mathcal{I} be the minimal model of P and \mathcal{I}' be the minimal model of $AD(P) \cup EDB(P)$.
- For every $p \in \mathcal{P}_{IDB}(P)$, $p \neq \text{answer}$, and every $\beta \in \text{valid}(p)$:

$$\mathcal{I}'[p-\beta] = \mathcal{I}[p].$$

- Furthermore: $\mathcal{I}'[\text{answer}] = \mathcal{I}[\text{answer}]$.

I.e. this part of the transformation has not changed the minimal model in an important way. It only renamed predicates (and possibly duplicated them).

Magic Set Transformation (1)

Notation:

- For a literal $A = p_{-}\beta(t_1, \dots, t_n)$ let

$$\text{magic}[A] := m_{-}p_{-}\beta(t_{i_1}, \dots, t_{i_k}),$$

where $1 \leq i_1 < \dots < i_k \leq n$ are the argument positions with $\beta(i_j) = b$.

- For a literal $A = \text{answer}(X_1, \dots, X_n)$ with the special predicate `answer` let $\text{magic}[A] := \text{true}$.

Example:

- $\text{magic}[\text{parent_bf}(X, Y)] := m_parent_bf(X).$

Magic Set Transformation (2)

Definition:

- Let $P = EDB(P) \cup IDB(P)$ be a logic program and $AD(P)$ be the version of $IDB(P)$ with explicit binding patterns.
- Then $MAG(P)$ contains the following rules:
 - For each rule $A \leftarrow B_1 \wedge \dots \wedge B_m$ in $AD(P)$:
$$A \leftarrow magic[A] \wedge B_1 \wedge \dots \wedge B_m.$$
 - For each rule $A \leftarrow B_1 \wedge \dots \wedge B_m$ in $AD(P)$ and every $i \in \{1, \dots, m\}$ with $pred(B_i) \in \mathcal{P}_{IDB}(P)$:
$$magic[B_i] \leftarrow magic[A] \wedge B_1 \wedge \dots \wedge B_{i-1}.$$

Magic Set Transformation (3)

Definition:

- Rules of the type

$$A \leftarrow \text{magic}[A] \wedge B_1 \wedge \cdots \wedge B_m$$

are called “**modified rules**”.

- Rules of the type

$$\text{magic}[B_i] \leftarrow \text{magic}[A] \wedge B_1 \wedge \cdots \wedge B_{i-1}$$

are called “**magic rules**”.

- A “**magic fact**” is a fact (ground atom) of the form $m_p_ \beta$.
- All other facts are called “**non-magic facts**”.

Magic Set Transformation (4)

Exercises:

- Check that the defined “magic set” transformation gives the result on Slide 8 and 9 (without suffix **bf**) for the “grandparent” example (Slide 4).
- Compute the result of the transformation for **append** with binding pattern **bbf**.

```
append([], L, L) ← true.  
append(X, L, Y) ← cons(H, T, X) ∧ cons(H, TL, Y) ∧  
                  append(T, L, TL).  
answer(X)        ← append([a], [b, c], X).
```

Correctness (1)

Lemma:

- Let \mathcal{I}_{AD} be the minimal model of $AD(P) \cup EDB(P)$, \mathcal{I}_{MAG} be the minimal model of $MAG(P) \cup EDB(P)$.
- For all non-magic facts A the following holds:

- If $\mathcal{I}_{MAG} \models A$, then $\mathcal{I}_{AD} \models A$.

Proof Sketch: Induction on the number of derivation steps. A non-magic fact can only be derived by a “modified rule”, but this is only a restricted version of the corresponding rule in $AD(P)$.

- If $\mathcal{I}_{AD} \models A$ and $\mathcal{I}_{MAG} \models magic[A]$, then $\mathcal{I}_{MAG} \models A$.

Proof Sketch: Induction on the number of derivation steps of A from \mathcal{I}_{AD} .

Correctness (2)

Theorem:

- Let \mathcal{I}_{AD} and \mathcal{I}_{MAG} be as above, and \mathcal{I} be the minimal model of $EDB(P) \cup IDB(P)$. Then the following holds:

$$\mathcal{I}_{MAG}[\text{answer}] = \mathcal{I}_{AD}[\text{answer}] = \mathcal{I}[\text{answer}].$$

- I.e. the transformed program is equivalent to the original program in the sense that it returns the same answers.

The two programs are not logically equivalent. Actually, that is not even defined, because the programs are based on different signatures. The programs could be called “answer-equivalent”.

Correctness (3)

Theorem:

- Let P be range-restricted with respect to $valid$.
- Then $MAG(P)$ is range-stricted with respect to $valid'$, where

$$valid'(q) := \begin{cases} \{\mathbf{f} \dots \mathbf{f}\} & \text{if } q \text{ has the form } p_{-\beta}/m_{-p_{-\beta}} \\ valid(q) & \text{otherwise.} \end{cases}$$

- I.e. the transformed program can be evaluated by iteration of the T_P -Operator.

Contents

- 1 Introduction
- 2 SIP-Strategies, Adorned Program
- 3 The Magic Set Transformation
- 4 Improvements

Supplementary Predicates (1)

Example:

- Rule from $AD(P)$, all B_i with IDB-predicates:

$$A \leftarrow B_1 \wedge B_2 \wedge B_3.$$

- Result of the magic set transformation:

$$magic[B_1] \leftarrow magic[A].$$

$$magic[B_2] \leftarrow magic[A] \wedge B_1.$$

$$magic[B_3] \leftarrow magic[A] \wedge B_1 \wedge B_2.$$

$$A \leftarrow magic[A] \wedge B_1 \wedge B_2 \wedge B_3.$$

- The same joins and selections are computed several times.

Supplementary Predicates (2)

Solution:

- Compute each join only once, store the result in a new “supplementary predicate”:

$$\begin{aligned} \textit{magic}[B_1] &\leftarrow \textit{magic}[A]. \\ S_1 &\leftarrow \textit{magic}[A] \wedge B_1. \\ \textit{magic}[B_2] &\leftarrow S_1. \\ S_2 &\leftarrow S_1 \wedge B_2. \\ \textit{magic}[B_3] &\leftarrow S_2. \\ A &\leftarrow S_2 \wedge B_3. \end{aligned}$$

- The arguments of the S_i are those variables from $\textit{magic}[A] \wedge B_1 \wedge \dots \wedge B_i$, that are still needed, i.e. that occur in B_{i+1}, \dots, B_m, A .

Supplementary Predicates (3)

- The deductive DBMS CORAL uses this method (“Magic Sets with Supplementary Predicates”).

CORAL is installed on our SUN machines. Before calling it, one must use `setenv CORALROOT /usr/central/coral`. Then it can be called with `$CORALROOT/bin/coral`. A manual is in `$CORALROOT/doc/manual.ps` (use `gv` to display it). Homepage: <http://www.cs.wisc.edu/coral/>.

In CORAL, rules are written into modules, an example `sg.P` is shown on the next slide. Facts are written outside modules, e.g. into `*.F`-files. The files are processed with, e.g., `consult(sg.P)`. For modules, this does the magic set transformation, the result is stored in `sg.P.M` (quite readable, i.e. one can look at the result of the transformation).

Strings are written, e.g., `"abc"`. If it has the form `[a-z][a-zA-Z0-9_]*`, no `"` is needed. Computation: `Y = X+1`. Query syntax: `? sg(julia, X)`.

To get all answers immediately: `clear(interactive_mode)`. Also useful: `help.`, `quit.`, `list_rels`.

Supplementary Predicates (4)

Example (CORAL):

- Rules in Coral are written in modules, with exported predicates and possible binding patterns defined.

```
module same_generation_cousins.
```

```
export sg(bf).
```

```
sg(X,X) :- person(X).
```

```
sg(X,Y) :- parent(X,Xp), sg(Xp,Yp), parent(Y,Yp).
```

```
end_module.
```

`export` also works with several binding patterns, e.g. `sg(bf,ff)`. More bound arguments in the call are possible, but not so efficient.

Supplementary Predicates (5)

- Example:

$sg(X, Y) \leftarrow parent(X, Xp) \wedge sg(Xp, Yp) \wedge parent(Y, Yp).$

- Result (if `parent` is an EDB-predicate):

$$\begin{aligned} sup_2_1(X, Xp) &\leftarrow m_sg_bf(X) \wedge \\ &\quad parent(X, Xp). \\ m_sg_bf(X) &\leftarrow sup_2_1(X, Xp). \\ sg_bf(X, Y) &\leftarrow sup_2_1(X, Xp) \wedge \\ &\quad sg_bf(Xp, Yp) \wedge \\ &\quad parent(Y, Yp). \end{aligned}$$

- In CORAL, the i -th supplementary predicate of the n -th rule is named “ sup_n_i ”.

Supplementary Predicates (6)

- Example (continued):

$$\text{sg}(X, X) \leftarrow \text{person}(X).$$

- The transformation of this non-recursive rule that only uses an EDB-predicate is easy:

$$\text{sg_bf}(X, X) \leftarrow \text{m_sg_bf}(X) \wedge \text{person}(X).$$

- Coral also supports other transformations.

Try one of: "@magic+", "@sup_magic+" (this is the default),
"@factoring+", "no_rewriting+", "sup_magic_indexing+".

- The SIP-strategy is left-to-right.

Supplementary Predicates (7)

Note:

- In this method, magic sets are always directly derived from supplementary predicates.
- One can try to replace the magic predicates by the supplementary predicates.
- If a magic predicate is defined by only one rule (only one call of p_{β}), this is simply a macro-expansion.
- Otherwise, one would have to duplicate rules.
- Depending on the application, it might be an advantage to distinguish different calls of a predicate.

Rectification (1)

- Problem/Example:

$$\begin{aligned}p(Y_1, Y_2, Y_3) &\leftarrow q(X, X, Y_1, Y_2, Y_3). \\q(a, b, Y_1, Y_2, Y_3) &\leftarrow r(Y_1) \wedge r(Y_2) \wedge r(Y_3).\end{aligned}$$

- The basic magic set method as explained above distinguishes only between “bound” and “free” argument positions.
- It calls q with the binding pattern $fffff$.
- Suppose that there are n facts about r . Then n^3 facts about q will be derived, although the rule does not match the call.

Rectification (2)

- Since the body of the first rule is not unifiable with the head of the second rule, SLD-resolution would immediately stop (without looking at the *r*-facts).

Such a situation probably does not happen often in practice. But since one wants to prove that magic sets are (in some sense) as efficient as (or really as goal-directed as) SLD-resolution, this is a problem.

- The magic set method can pass only concrete values for the arguments to a called predicate.
- The basic method cannot pass the information to *q* that the first two arguments must be equal.

Rectification (3)

Definition:

- A logic program is rectified iff no body literal contains the same variable more than once, i.e. for every body literal $p(t_1, \dots, t_n)$, if $t_i = t_j$ for $i \neq j$, then t_i is a constant.

Remark:

- In the magic predicates, free argument positions are projected away (not represented). If the program is rectified, this does not lead to a loss of information.

Rectification (4)

- Every logic program that does not contain function symbols (structured terms) can be transformed into an equivalent, rectified program.

Again, equivalent means that it produces the same answer.

- The rectification is done by introducing predicate variants that contain at different argument positions the same arguments: $p^{(i_1, \dots, i_n)}(t_1, \dots, t_k)$ corresponds to $p(t_{i_1}, \dots, t_{i_n})$, e.g.
 - $q^{(1,1,2,3,4)}(X, Y_1, Y_2, Y_3)$ means $q(X, X, Y_1, Y_2, Y_3)$.

Rectification (5)

- One specializes the rules of the original predicate once for each such predicate variant:
 - E.g., the critical rule $q(a, b, \dots) \leftarrow \dots$ is deleted in the specialization for the predicate variant $q^{(1,1,2,3,4)}$.

Note that this is a predicate name. In practice, one of course has to encode it without exponent, parentheses, commas.
- In general, one tries to unify the rule head with $p(t_{i_1}, \dots, t_{i_n})$. If it is unifiable, the result is encoded with the new predicates: $p^{(i_1, \dots, i_n)}$ in the head, and in the body as needed to ensure rectification.

Rectification (6)

- As explained above, rectification is applied before the adorned program is computed and the SIP-strategy is applied.
- However, one could do the rectification also together with the adornment.
- Then binding patterns consist not of **b** and **f**, but of “constant”, “variable-1”, “variable-2”, and so on.

Note that it is no problem if a bound variable appears twice in a body literal. So one might need fewer predicate variants in this way. Note also that if one wants to come close to SLD-resolution, only SIP-strategies are interesting that do not ignore existing bindings.