Logic Programming and Deductive Databases

Chapter 7: Practical Prolog Programming

Prof. Dr. Stefan Brass
Martin-Luther-Universität Halle-Wittenberg

Summer 2023

http://www.informatik.uni-halle.de/~brass/lp23/

Objectives

After completing this chapter, you should be able to:

- explain the effect of the cut.
- enumerate some applications of the cut.
- explain why the cut is somewhat dangerous.

  Give an example of an unexpected behaviour.
- decide whether a recursive predicate is tail recursive.
- explain how to translate loops from imperative programs to Prolog code.
- explain different ways to represent arrays in Prolog.
- write Prolog programs for practical applications.

Contents

The Cut: Effect (1)

- The cut, written "!" in Prolog, removes alternatives that otherwise would have been tried during backtracking. E.g. consider this rule:

  $$p(t_1, \ldots, t_k) \;\texttt{:-}\; A_1,\; \ldots,\; A_m,\; !,\; B_1,\; \ldots,\; B_n.$$

- Until the cut is executed, processing is as usual.
- When the cut is reached, all previous alternatives for this call to the predicate p are removed:
  - No other rule about p will be tried.
  - No other solutions to the literals $A, \ldots, A_m$ will be considered.

The Cut: Effect (2)

- Example:

```
p(X) :- q(X), !, r(X).
p(X) :- s(X).
q(a).
q(b).
r(X).
s(c).
```

- With the cut, the query `p(X)` returns only `X=a`.
- Without the cut, the solutions are `X=a`, `X=b`, `X=c`.
- Exercise: Can the second rule about `p` ever be used?

The Cut: Effect (3)

Four-Port Model without Cut:

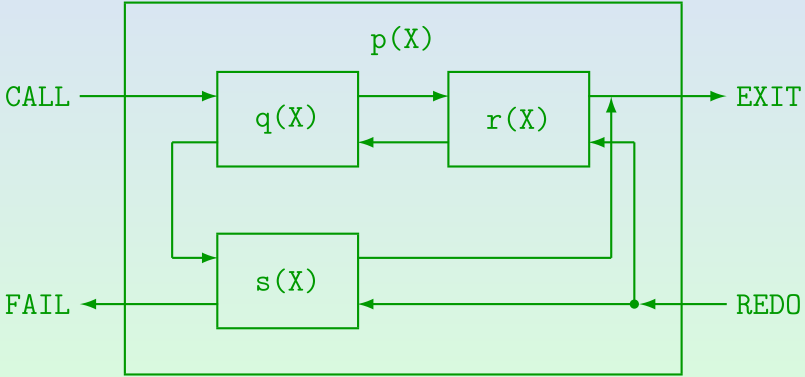

The Cut: Effect (4)

Four-Port Model with Cut:

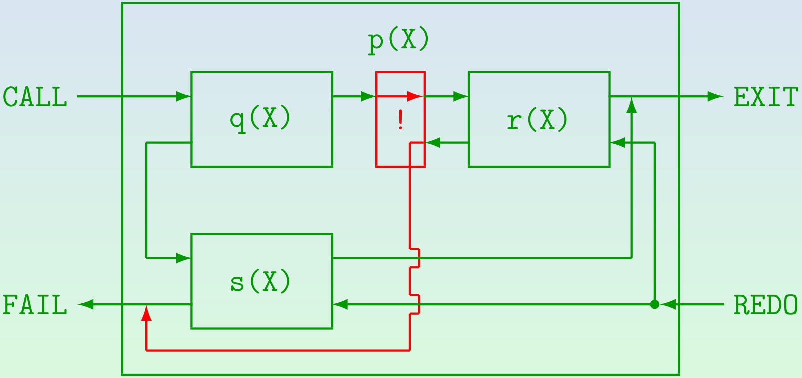

The Cut: Effect (5)

- A call to the cut immediately succeeds (like `true`).
- Any try to redo the cut not only fails, but immediately fails the entire predicate call.
- In the SLD-tree, the cut "cuts away" all still open branches between
  - the node where the cut was introduced (i.e. the child of which contains the cut), and
  - the node where the cut is the selected literal.

The Cut: Effect (6)

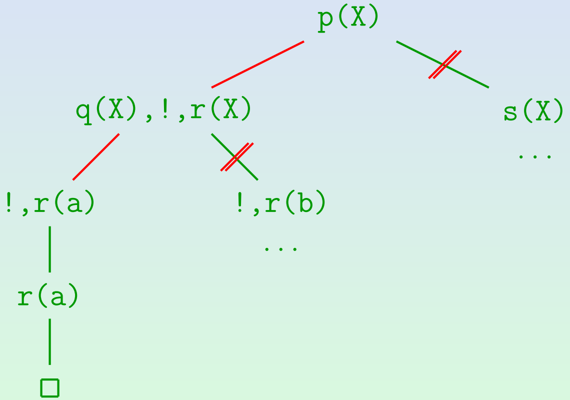

The Cut: Effect (7)

- Before and after the cut, the backtracking is normal (only not through the cut):

```
p(X,Y) :- q1(X), q2(X,Y).
p(X,Y) :- r1(X), r2(X), !, r3(X,Y), r4(Y).
p(X,Y) :- s(X,Y).
q1(a).
q1(b).     q2(b,c).
r1(d).
r1(e).     r2(e).
r1(f).     r2(f).
r3(e,g).
r3(e,h).   r4(h).
r3(f,i).   r4(i).
s(j,k).
```

- The query `p(X,Y)` has solutions `X=b,Y=c` and `X=e,Y=h`.

Contents

Cut: Improving Runtime (1)

- One application of the cut is to improve the runtime of a program by eliminating parts of the proof tree that cannot yield solutions or at least cannot yield any new solutions.
- Consider the predicate `abs` that computes the absolute value of a number:

```
abs(X,X) :- X >= 0.
abs(X,Y) :- X =< 0, Y is -X.
```

- When the first rule is successful, it is clear that the second rule does not have to be tried.

Cut: Improving Runtime (2)

- Consider now the goal

  `p(X), abs(X,Y), Y > 5`

  with the facts `p(3)`, `p(0)`, `p(-7)`.

- First `p(X)` succeeds with `X` bound to `3`, then `abs(3,Y)` succeeds for `Y=3`, but then `3 > 5` fails.

- Now backtracking would normally first try to find an alternative solution for `abs(3,Y)`, since there is another rule about `abs` that has not yet been tried.

- This is obviously useless, and the runtime can be improved by immediately backtracking to `p(X)`.

Cut: Improving Runtime (3)

- With the cut, one can tell the Prolog system that when the first rule succeeds, the second rule cannot give anything new:

```
abs(X,X) :- X >= 0, !.
abs(X,Y) :- X =< 0, Y is -X.
```

- Of course, one could have (should have) written the condition in the second rule `X < 0`.
- Then some (but not all) Prolog systems are able to discover themselves that the rules are mutually exclusive.

Cut: Improving Space (1)

- Making clear that a predicate has no other solution improves also the space (memory) efficiency.
- The Prolog system must keep a record ("choicepoint") for each predicate call that is not yet complete (for backtracking into the predicate call later).
- Even worse, certain data structures within the Prolog system must be "frozen" when it is necessary to support later backtracking to this state.
- Then e.g. variable bindings must be logged (on the "trail") so that they can later be undone.

Cut: Improving Space (2)

- In imperative languages, when a procedure call returns, its stack frame (containing local variables and other information) can be reused.
- In Prolog, this is not always the case, because it might be necessary to reactivate the procedure call and search for another solution.
- E.g. consider the following program:

```
p(X) :- q(X), r(X).
q(X) :- s(X), t(X).
s(a). s(b). t(a). t(b). r(b).
```

Cut: Improving Space (3)

- The call `q(X)` first exits with `X=a`, but then `r(a)` fails, thus the call `q(X)` is entered again, which in turn reactives `s(X)`.

  Upon backtracking, also the binding of `X` must be undone.

- In the above example, not much can be improved, because there really are alternative solutions.
- However, when a predicate call has only one solution, it should be executed like a procedure call in an imperative language.

Cut: Improving Space (4)

- Predicate calls that can have at most one solution are called deterministic.

  Sometimes one calls the predicate itself deterministic, but then one usually has a specific binding pattern in mind. E.g. `append` is deterministic for the binding pattern `bbf`, but it is not deterministic for `ffb`.

- For efficient execution, it is important that the Prolog system understands that a predicate call is deterministic. Here a cut can help.

  Actually, the cut in the definition of `abs` makes the predicate deterministic. In general, it might be important that `abs(0,X)` succeeds "two times", Prolog is not allowed to automatically remove one solution. Deductive databases are set-oriented, there more powerful optimizers are possible.

Cut: Improving Space (5)

- Consider abs applied to a list:

```
abs_list([], []).
abs_list([X|R], [Y|S]) :- abs(X, Y),
                          abs_list(R, S).
```

- When the Prolog system thinks that abs is nondeterministic, it will keep the stackframe for each call to abs (and for the calls to abs_list).

- When a predicate calls a nondeterministic predicate, it automatically becomes nondeterministic, too.

  Only for the last body literal of the last rule about a predicate, the stack frame of the predicate is reused (under certain conditions), and thus does not remain, even when this body literal is non-deterministic.

Cut: Improving Space (6)

- In the above example, making abs deterministic (by means of a cut) is a big improvement.
- Then most Prolog systems will automatically deduce that also abs_list is deterministic.

  For the only possible binding patterns bf and bb.
- Usually, the outermost functor of the first argument is considered: Since it is "[]" for the first rule, and "." for the second, always only one of the two rules is applicable (if the first argument is bound).

Cut: Improving Space (7)

- It is also possible to remove unnecessary stack frames at a later point.
- E.g. suppose that `abs` (and thus `abs_list`) remain nondeterministic, and consider the goal:

```
abs_list([-3,7,-4], X), !, p(X).
```

- The call to `abs_list` will leave many stack frames behind, but these are deleted by the cut.

  It is probably better style to avoid the nondeterminism at the place where it occurs. However, one should not use too many cuts, and it might be easier to clean up the stack only at a few places.

Contents

Cut: If-Then-Else (1)

- The cut is also used to encode an “if then else”.
- Consider the following predicate:

```
p(X, Y) :- q1(X), !, r1(X, Y).
p(X, Y) :- q2(X), !, r2(X, Y).
p(X, Y) :- r3(X, Y).
```

- This is equivalent to (assuming that q1 and q2 are deterministic):

```
p(X, Y) :- q1(X), r1(X, Y).
p(X, Y) :- \+ q1(X), q2(X), r2(X, Y).
p(X, Y) :- \+ q1(X), \+ q2(X), r3(X, Y).
```

Cut: If-Then-Else (2)

- The formulation with the cut is a bit shorter.

  The difference becomes the bigger, the more cases there are.

- Furthermore, the runtime is shorter: In the version without the cut, `q1(X)` is computed up to three times.

- But removing the cut in first version would completely change the semantics of the program.

  The cut is no longer only an "optimizer hint".

Cut: If-Then-Else (3)

- The logical semantics of programs with negation as failure ("\+") has be extensively studied and there are good proposals.
- I do not know of successful tries to give the cut a clear logical (declarative) semantics.

  The cut can basically be understood only operationally. One problem is that the cut is used for many different purposes, and it might be difficult to automatically discover for which one.

- Pure logic programmers try to avoid the cut, at least when it affects the logic of the program.

Cut: If-Then-Else (4)

- Prolog has an "if-then" operator `->` that can be used to have the advantages of the cut, while making the logical intention clear.
- E.g. one could write the above procedure as

```
p(X, Y) :- q1(X) -> r1(X,Y);
           q2(X) -> r2(X,Y);
           r3(X,Y).
```

- *A* `->` *B* has basically the same effect as *A*`, !, `*B*.

  However, if there should be further rules about `p`, this cut does not remove the possibility to try these rules. It does remove alternative solutions for *A*, and it does remove the possibility to try the disjunctive alternatives within the rule.

Cut: Negation

- Conversely, one can implement negation as failure with the cut (`not` is only another name for `\+`):

```
not(A) :- call(A), !, fail.
not(_).
```

- The first rule ensures that if `A` succeeds, `not(A)` fails.
- The second rule makes `not(A)` true in all other cases (i.e. when `A` fails).

  Of course, if `A` should run into an infinite loop, also `not(A)` does not terminate.

Cut: One Solution (1)

- Suppose that email addresses of professors are stored as facts, and that the same person can have several email addresses:

```
prof_email(brass, 'sbrass@sis.pitt.edu').
prof_email(brass, 'brass@acm.org').
prof_email(spring, 'mspring@sis.pitt.edu').
...
```

- The cut can be used to select a single address of a given professor:

```
prof_email(brass, E), !, send_email(E).
```

Cut: One Solution (2)

- Prolog has a built-in predicate `once` that can be used instead of the cut:

```
once(prof_email(brass, E)), send_email(E).
```

- `once` is defined as:

```
once(A) :- call(A), !.
```

- In the example, the following is equivalent:

```
prof_email(brass, E) -> send_email(E).
```

  However, the solution with `once` makes the intention clearer.

Contents

Cut: Dangers (1)

- The cut can make programs wrong if predicates are called with unexpected binding patterns.
- E.g. the predicate for the absolute value can also be written as follows (using the cut as in the if-then-else pattern):

```
abs(X,X) :- X >=0, !.
abs(X,Y) :- Y is -X.
```

- Since the second rule is executed only when the first rule is not applicable, it might seem that the test `X =< 0` used earlier is superfluous.

Cut: Dangers (2)

- This is indeed true for the binding pattern `bf`, but consider now the call `abs(3,-3)`!
- In general, the rule is that the cut must be exactly at the point where it is clear that this is the right rule: Not too early and not too late.
- Here the unification must happen after the cut:

```
abs(X,Y) :- X >= 0, !, X = Y.
abs(X,Y) :- Y is -X.
```

- This would work also with binding pattern `bb`.

Cut: Dangers (3)

- Consider this predicate:

```
person(X, male)   :- man(X), !.
person(X, female) :- woman(X).
```

- Since `man` and `woman` are disjoint, the cut was only added to improve the efficiency.
- It works if `person` is called with binding pattern `bf` or `bb`. However, consider what happens if `person` is called with binding pattern `ff`!

  It is interesting that here the more general binding pattern poses a problem, whereas in the `abs` example, the more specific binding pattern is not handled.

Types of Cuts

- Cuts in Prolog programs are usually classified into
  - Green Cuts: Do not modify the logical meaning of the program, only improve the runtime/space efficiency.

    Some authors also distinguish blue cuts: In this case, a good Prolog system should be able to determine itself that there are no further solutions. Blue cuts are intended only for very simple Prolog systems. “Grue Cuts”: Green or blue cuts.

  - Red cuts: Modify the declarative meaning of the program.

    Good Prolog programmers try to use red cuts only very seldom.

Cut: Summary, Outlook

- The cut is necessary for efficient Prolog programming, but it destroys the declarative meaning of the programs and can have unexpected consequences.
- The better Prolog implementations get, the less important will be the cut.
- Newer logic programming languages usually try to replace the cut by other constructs that have a more declarative meaning.
- If possible, use `->`, `\+`, `once` instead.
- Use the cut only as last resort.

Contents

Prolog vs. Pascal

- "Prolog is different, but not that different."

  This citation is probably from O'Keefe, The Craft of Prolog.

- In general, one can translate a given imperative algorithm (from Pascal, C, etc.) into Prolog.

  The resulting program might be not the best possible program for the task, just as a word-by-word translation from e.g. German to English gives bad English. But at least, if one knows how to solve a problem in an imperative language, one should also be able to write a Prolog program for it.

- The goal of this section is to teach some typical patterns of Prolog programming.

Data Types (1)

Pascal	Prolog
integer	integer
real	float
char	ASCII-code (integer) atom
string	list of ASCII-codes atom string (in some Prologs)
file	stream atom (alias, in some Prologs) switching standard IO

Data Types (2)

Pascal	Prolog
enumeration type	set of atoms
(variant) record (union/struct in C)	composed term `functor(field1, ..., fieldN)`
array	list set of facts: `a(i, valI)` term: `a(val1, ..., valN)`
pointer	structured terms (e.g., lists) otherwise like array index
—	partial data structures (terms with variables)

Variables (1)

- One can assign a value to a Prolog variable only once.

  This is the biggest difference to imperative programming.

- Afterwards it is automatically replaced everywhere by its value.

  I.e. it ceases to exist as a variable.

- Thus, there is no possibility to assign a different value during normal, forward execution.

  Of course, with backtracking, one can go back to the point where the variable was still unbound. But then all other variable bindings that happend since that point in time are also undone.

Variables (2)

- Thus, one uses a different variable for every value:
  - ```
    procedure p(n: integer, var m: integer);
    begin
        m := n * 2;
        m := m + 5;
        m := m * m
    end
    ```
  - ```
    p(N, M) :-
        M1  is  N * 2,
        M2  is  M1 + 5,
        M   is  M2 * M2.
    ```

    This is an artificial example. Normally, one would compute the value of m in a single expression.

Variables (3)

- Using a new variable for every assignment is obviously possible for sequential/linear code.
- Loops are formulated in Prolog by recursion, thus one can also get a fresh set of variables for every iteration (see Slide 57 for more efficient solution):
  - **for** i := 1 **to** n **do** writeln(i);
  -
    ```
    loop(N)          :- loop_body(1, N).
    loop_body(I, N)  :- I > N.
    loop_body(I, N)  :- I =< N, write(I), nl,
                        Next_I is I + 1,
                        loop_body(Next_I, N).
    ```

Variables (4)

- For variables passed between procedures (in/out-parameters), a Pascal variable is split into two Prolog variables: One for the input value, and one for the output value (“accumulator pair”).

  - ```
    procedure double(var n: integer);
    begin
        n := n * 2
    end
    ```

  - ```
    double(N_In, N_Out) :-
        N_Out  is  N_In * 2.
    ```

Variables (5)

- Global variables in Pascal should be made predicate parameters in Prolog.

  Values that are passed unchanged from predicate to predicate are called "context arguments".

- If there are too many global variables, one can pack them into a structure (composed term) which can be passed as a unit. An example is shown on the next slide.

  Now every predicate that needs access to the "global state" has two additional parameters `+StateIn` and `-StateOut`. If the predicate does not need to change the state, a single parameter `+State` suffices. The state can encapsulate any number of components, and one can easily add components if the main code uses only `getA(+State, -A)` and `setA(+StateIn, +A, -StateOut)` predicates for accessing the state (see next slide). One can also check the validity of the component values in the `setA`-predicates.

Variables (6)

Data Structure with Three Components X, Y, Z:

- The variable values are packed into a term `xyz(X,Y,Z)`.
- `init(-Obj)` returns object with initial values:

```
init(xyz(0,0,0)).
```

- `getA(+Obj, -A)` returns value of component `A`:

```
getX(xyz(X,_,_), X).
getY(xyz(_,Y,_), Y).
getZ(xyz(_,_,Z), Z).
```

- `setA(+In, +A, -Out)` changes value of component `A`:

```
setX(xyz(_,Y,Z), X, xyz(X,Y,Z)).
setY(xyz(X,_,Z), Y, xyz(X,Y,Z)).
setZ(xyz(X,Y,_), Z, xyz(X,Y,Z)).
```

Variables (7)

- One can represent a global variable also with a fact in the dynamic database:
  - x := x+1
  - x(X), retract(x(X)), !, X1 is X+1, assert(x(X1)).

    There is always only one fact about the predicate `x` which contains the current value of `X`.

- Some Prolog systems have support for destructive assignments and global variables, but that is very system-dependent.

  In GNU-Prolog, there is, e.g., `g_assign/2`, `g_read/2`, `g_array_size/2`. In SWI-Prolog, see `flag/3`, `setarg/3`. In Sepia/ECLiPSe, see `setval/2`.

Conditions: If, Case (1)

- Example:

  procedure min(i, j: **integer**; **var** m: **integer**);
  begin if i < j **then** m := i **else** m := j **end**

- There are basically three possibilities to translate conditional statements:
  - One rule per case, body contains the complete condition (i.e. the negation of all previous cases):

```
min(I, J, M)  :-  I < J,        M = I.
min(I, J, M)  :-  \+ (I < J),   M = J.
```

Of course, one would write I >= J instead of \+ (I < J). But not all conditions can be inverted so simply.

Conditions: If, Case (2)

- Translations of **if**-statements, continued:
  - A rule per case, each starts with the "else if"-condition and a cut:

```
min(I, J, M)  :-  I < J, !,  M = I.
min(I, J, M)  :-             M = J.
```

  - Using the conditional operators, i.e. a rule with a body of the form (Cond-> Then ; Else):

```
min(I, J, M)  :-  (I < J -> M = I; M = J).
```

  - If possible, it is best to express the condition in the head of the rule (see next slide).

Conditions: If, Case (3)

- Example (condition in the rule head):
  - ```
    procedure p(c: color; var i: integer);
    begin
        case c of red:    i := 1;
                  green:  i := 2;
                  blue:   i := 3;
        end
    end
    ```
  - ```
    p(red, I)    :-  I = 1.
    p(green, I)  :-  I = 2.
    p(blue, I)   :-  I = 3.
    ```

    Of course, one would simplify this further to, e.g., p(red, 1).

Conditions: If, Case (4)

- The first solution (complete condition in every rule) is logically cleanest, it also permits to understand each rule in isolation.
- However:
  - There is a certain duplication of code.
  - In many Prolog systems, it will leave a choice point behind if the first alternative is chosen.

    Some Prolog systems are intelligent enough to understand that `I < J` and `I >= J` exclude each other, thus no choice point is needed. The solution with the cut or `->` does not have this problem.

Conditions: If, Case (5)

- In general, it is better to use the explicit conditional operator -> instead of the cut !, because this makes the purpose of the cut clear.
  - However, this makes the syntactical structure of the rules more complicated.
- Most Prolog systems have an index (hash table) over the outermost functor of the first argument.
  - Thus, if it is possible to code the condition in the rule head (first argument), this will be especially efficient, and no choicepoint will be generated.

Loops: Tail-Recursion (1)

- Loops are usually written as end-recursions.
- One should try to make sure that only a constant amount of memory is needed, not an amount linear in the number of executions of the loop body.
- A Prolog system reuses the memory of a rule invocation when the last body literal is called and there are no more alternatives.

  If necessary, one can use a cut to make clear that other rules are not applicable. It is best when the recursive call is the last literal of the last rule about the predicate.

Loops: Tail-Recursion (2)

- Example (list length):

```
procedure length(l: list; var n: integer);
begin
    n := 0;
    while l <> nil do begin
        n := n + 1;
        l := l↑.next
    end
end
```

Loops: Tail-Recursion (3)

- Direct Translation to Prolog:

```
length(L, N) :-
    length(L, 0, N).

length(L, N_In, N_Out) :-
    L = [ ], !,
    N_Out = N_In.

length(L, N_In, N_Out) :-
    N_Next is N_In + 1,
    [_|L_Next] = L,
    length(L_Next, N_Next, N_Out).
```

Loops: Tail-Recursion (4)

- Using the rule head for the conditions (makes the cut unnecessary):

```
length(L, N) :-
    length(L, 0, N).

length([ ], N_In, N_Out) :-
    N_Out = N_In.

length([_|L_Next], N_In, N_Out) :-
    N_Next is N_In + 1,
    length(L_Next, N_Next, N_Out).
```

Loops: Tail-Recursion (5)

- Alternative (elegant, but not tail recursive):

```
length([ ], 0).
length([_|L_Rest], N) :-
        length(L_Rest, N_Rest),
        N is N_Rest + 1.
```

- In this case, the system must return to a rule invocation after the recursive call.
- Thus, many systems will need memory that is linear in the length of the list. But efficiency is not all!

  The memory will become free after the call to the length predicate. If the lists are not extremely long, the more elegant solution should be preferred.

Loops: Tail-Recursion (6)

- Consider again the for-loop:
  - **for** i := 1 **to** n **do** writeln(i);
  -
```
loop(N)            :-  loop_body(1, N).
loop_body(I, N)    :-  I > N.
loop_body(I, N)    :-  I =< N, write(I), nl,
                       Next_I is I + 1,
                       loop_body(Next_I, N).
```

- The last call (I = N+1) probably leaves a choice point behind.

  A good Prolog system might discover that the two conditions I > N and I =< N are mutually exclusive. Then there will be no choicepoint.

- However, for all other calls, tail recursion works: The memory complexity is still constant.

Loops: Tail-Recursion (7)

- The stack frame of a rule invocation

$$p(\dots) \leftarrow B_1 \wedge \dots \wedge B_n$$

is recycled for the last call in the rule (B_n) if

  - this is the last rule about p

    Or the compiler knows that following rules are not applicable.

  - and no alternatives remain for the previous body literals $B_1, \dots, B_{n-1}$.

    Thus, it is certainly not necessary to enter this rule invocation again. If necessary and logical possible, one might add a cut before the last literal of the last rule to eliminate choice points generated by $B_1, \dots, B_{n-1}$.

- It does not prevent this optimization if the last literal B_n might have several solutions.

Loops: Tail-Recursion (8)

- In the example, one could use a cut to make clear that the second rule is no alternative when the first rule is applicable:

```
loop_body(I, N)   :-   I > N, !.
loop_body(I, N)   :-   write(I), nl,
                       Next_I is I + 1,
                       loop_body(Next_I, N).
```

- However, too many cuts are bad style.
- When testing an auxiliary predicate in isolation, it is nice if the Prolog system does not wait for the user to press "`;`", only to answer `false`/`no` afterwards.

Loops: Backtracking (1)

- Some loops can also be written with "repeat" and backtracking.
- This is only possible if only the side effect of the loop body is important (input/output, changes to the dynamic DB), but no variable bindings must be kept from one iteration to the next.
- If it is possible, it is very efficient.

  Not only the stack frames of predicate invocations are reused, but also all term structures that were built up. This is important for simple Prolog systems that have no garbage collection: There, space on the heap is recycled only upon backtracking. For loops that run an indefinite amout of time (command loops etc.), this should be used.

Loops: Backtracking (2)

- ```
  procedure skip(c: character);
  var c1: character;
  begin
      repeat
          read(c1)
      until c1 = c
  end
  ```

- ```
  skip(C) :- repeat,
             get_code(C1),
             C1 = C,
             !.
  ```

  The cut is important to avoid the re-execution of the loop if later in the program something fails and backtracking starts.

Loops: Backtracking (3)

- Another typical pattern for a backtracking loop is to iterate over all solutions to a predicate:

```
print_all_composers :-
        composer(FirstName, LastName),
        write(LastName),
        write(', '),
        write(FirstName),
        nl,
        fail.
print_all_composers.
```

The fact at the end ensures that the predicate ultimatively succees.