

# Incremental answer completion in the SLG-WAM

Terrance Swift, Alexandre Miguel Pinto, and Luís Moniz Pereira \*

Centro de Inteligência Artificial, Universidade Nova de Lisboa

**Abstract.** The SLG-WAM of XSB Prolog soundly implements the Well-Founded Semantics (WFS) for logic programs, but in a few pathological cases its engine treats atoms as undefined that are true or false in WFS. The reason for this is that the XSB does not implement the SLG ANSWER COMPLETION operation in its engine, the SLG-WAM – rather ANSWER COMPLETION must be performed by post-processing the table. This engine-level omission has not proven significant for applications so far, but the need for ANSWER COMPLETION is becoming important as XSB is more often used to produce well-founded residues of highly non-stratified programs. However, due to its complexity, care must be taken when adding ANSWER COMPLETION to an engine. In the worst case, the cost of each ANSWER COMPLETION operation is proportional to the size of a program  $P$ , so that the operation must be invoked as rarely as possible, and when invoked the operation must traverse as small a fragment as possible of  $P$ . We examine the complexity of ANSWER COMPLETION; and then describe its implementation and performance in XSB’s SLG-WAM such that the invocations of the operation are restricted, and which is limited in scope to Strongly Connected Components within a tabled evaluation’s Subgoal Dependency Graph.

Designers of logic programming engines must weigh the usefulness of operations against the burden of complexity they require. Perhaps the best known example is the *occurs check* in unification. Prologs derived from the WAM do not perform occurs check between two terms, since its cost may be exponential in the size of the terms. Rather, the occurs check must be explicitly invoked through the ISO predicate `unify_with_occurs_check/2` or a similar mechanism. For evaluating normal programs using tabling, checking for certain positive loops involves similar considerations. While most positive loops can be efficiently checked, positive subloops within larger negative loops are more difficult to detect, and account for the complexity of evaluating a program  $P$  according to WFS, which is  $atoms(P) \times size(P)$ , where  $atoms(P)$  is the number of atoms of  $P$  and  $size(P)$  is the number of rules of  $P$ .

As implemented in XSB, the SLG-WAM detects positive loops between tabled subgoals so that answers are not added to a table unless they are true, or are involved in a loop through negation and so are undefined at the time of their addition (termed *conditional answers*). As shown in Theorem 1 below, this sort of evaluation can be done in time linear in  $size(P)$ . However, a situation can arise where certain conditional answers are later determined to be true or false. This determination may break a negative loop, which uncovers a positive loop and makes the answers false. Within SLG, this situation is addressed by the ANSWER COMPLETION operation, which is not implemented within the currently available version of the SLG-WAM. So far, the lack of ANSWER

---

\* tswift@cs.sunysb.edu, amp@di.fct.unl.pt, lmp@di.fct.unl.pt

COMPLETION has not proven a problem for most programs. However, the SLG-WAM is increasingly being used to produce well-founded residues for highly non-stratified programs for applications involving intelligent agents (e.g. [2]), where the need for ANSWER COMPLETION is greater.

This paper examines issues involved in adding ANSWER COMPLETION to the SLG-WAM. We illustrate the situation of a positive loop begin uncovered when a negative loop is resolved through a concrete example, and then we provide a formal result on the contribution ANSWER COMPLETION makes to the complexity of computing WFS. We introduce an algorithm for efficiently performing ANSWER COMPLETION (subject to its complexity), and discuss performance results obtained by implementing it in the SLG-WAM. Due to space requirements, we must assume knowledge of tabled evaluation of WFS through SLG resolution [1] as well as certain data structures of the SLG-WAM [3].

*Example 1.* The following program is soundly, but not completely, evaluated by the SLG-WAM, where `tnot/1` indicates that tabled negation is used:

```
:- table p/1,r/0,s/0.
p(X):- tnot(s).      p(X):- p(X).
      s:- tnot(r).      s:- p(X).      r:- tnot(s),r.
```

The well founded model for this program has true atoms  $\{s\}$  and false atoms  $\{r, p(X)\}$ . Recall that literals that do not have a proof and that are involved in loops over default negation are considered *undefined* in WFS. Unproved literals involved only in positive loops, i.e., without negations, are *unsupported* and, hence, *false* in WFS. Accordingly,  $p(X)$ , whose second clause fails, is *false*; however, a query to  $p(X)$  in XSB indicates that  $p(X)$  is *undefined*. The reason is that during evaluation the engine detects a strongly connected component (SCC) of mutually dependent goals containing  $p(X)$ ,  $r$  and  $s$ , along with negative dependencies, and no answers for any of these goals. In such a situation, the SLG-WAM delays negative literals and continues execution. Here, the literal `tnot(s)` in the rule  $p(X) :- \text{tnot}(s)$  is delayed, producing an answer  $p(X) :- \text{tnot}(s) |$ , indicating that  $p(X)$  is conditional on a *delay list*, here `tnot(s)`. That answer is returned to the goal  $p(X)$  in the clause  $p(X) :- p(X)$  and a conditional answer  $p(X) :- p(X) |$  is derived. Later, a positive loop is detected for  $r$ , causing its truth value to become *false*. The failure of  $r$  causes  $s$  to become *true*, and SIMPLIFICATION removes the answer  $p(X) :- \text{tnot}(s) |$ . At this stage, however, no further simplification is possible for  $p(X) :- p(X) |$ , which is now unsupported.

The ANSWER COMPLETION operation addresses such cases by detecting positive loops in dependencies among conditional answers. More precisely, ANSWER COMPLETION marks *false* sets of answers that are not *supported*: i.e. conditional answers for completed subgoals that contain only positive, and no negative dependencies in their delay lists. The creation of unsupported answers are uncommon in the SLG-WAM because its evaluation is *delay minimal* – that is, the engine performs no unnecessary DELAYING operations [4]. Delay minimality reduces the need for simplification of dependencies among answers, and thereby the chances of uncovering positive loops among answers, as with the answer  $p(X) :- p(X) |$  above.

## 1. Complexity

We begin by showing that queries to programs that do not need ANSWER COMPLETION can be evaluated in  $\mathcal{O}(\text{size}(P))$ . Such programs include stratified ones, and also non-stratified programs that contain no positive loops within negative SCCs in their dynamic dependency graphs<sup>1</sup>.

**Theorem 1.** *Let  $Q$  be a query to a finite ground normal program  $P$ . Under a cost model with constant time access to all subgoals, nodes, and delay elements of each SLG forest in an evaluation, and constant time access to each clause in  $P$ , a partial SLG evaluation that does not perform ANSWER COMPLETION can be constructed that is linear in the size of  $P$ .*

The algorithm ITERATE ANSWER COMPLETION below iteratively applies ANSWER COMPLETION operations, calling Check Supported Answers() to perform a check for positive loops. Check Supported Answers() is an adaptation of Tarjan’s algorithm for SCC detection (cf. [http://en.wikipedia.org/wiki/Tarjan's\\_strongly\\_connected\\_components\\_algorithm](http://en.wikipedia.org/wiki/Tarjan's_strongly_connected_components_algorithm)), which is linear in  $\text{size}(P)$ . Note that in the worst case, ANSWER COMPLETION operations iteratively need to be applied, and that each time it is applied, a single atom would be found *false*. In that case, program evaluation would have a cost proportional to  $\text{atoms}(P) \times \text{size}(P)$ , which is equivalent to the known complexity for WFS.

## 2. Implementation of ANSWER COMPLETION

Within an SLG evaluation, a tabled subgoal can be marked as *complete*, which indicates that all possible answers have been produced for the subgoal, although SIMPLIFICATION and ANSWER COMPLETION operations may remain to simplify or delete conditional answers. Completed subgoals do not require execution stack space, but only table space, so that completing subgoals as early as possible is a critical step for engine efficiency. Accordingly the SLG-WAM performs *incremental completion* via a `completion` instruction, which maintains information about mutually dependent sets of subgoals (SCCs), and completes each SCC when all applicable operations have been performed. In addition to marking each subgoal  $S$  in an SCC as complete, if  $S$  failed (has no answers) the `completion` instruction may initiate SIMPLIFICATION for conditional answers that depend negatively on  $S$ . In terms of ANSWER COMPLETION, observe that any positive loop among the delayed literals of conditional answers must be contained within the SCC being completed, as each delayed literal was a selected literal before it was delayed. This incremental approach has several benefits. Performing ANSWER COMPLETION operation within the `completion` instruction restricts the space that any such operation needs to search. In addition, performing ANSWER COMPLETION after all other SIMPLIFICATION operations have been performed on answers within the SCC similarly reduces search space. As a final optimization, ANSWER COMPLETION is not required unless delaying has been performed within the SCC, a fact that is easily maintained using data structures in the SLG-WAM’s *Completion Stack*, which maintains information about SCCs. The pseudo code for `Iterate Answer Completion()`, which traverses all subgoals in the SCC using the *Completion Stack*, and checks each

---

<sup>1</sup> The proof of Theorem 1 is contained in an appendix of a fuller version of this paper available on request.

answer for support, deleting unsupported answers from the table and invoking SIMPLIFICATION operations, is presented in Figure 1. SIMPLIFICATION may remove further negative loops, and uncover new unsupported other answers as a side-effect. In such case, the ANSWER COMPLETION procedure should be executed once more, and this is guaranteed by the use of the *reached\_fixed\_point* flag. A fixed-point is reached when all answers within the scope of the SCC are known to be supported.

Algorithm Iterate Answer Completion(*CompletionStack*)

```

reached_fixed_point = FALSE;
while not reached_fixed_point
  reached_fixed_point = TRUE;
  foreach subgoal S in the Completion Stack
    foreach answer A for subgoal S
      if not Check_Supported_Answer(A)          /* A is unsupported */
        reached_fixed_point = FALSE;
        delete A;
        propagate A's deletion's simplifications;

```

**Fig. 1.** Algorithm ITERATE ANSWER COMPLETION.

**Check Supported Answer** This procedure (Figure 2) does the actual check of whether a (positive) answer is unsupported. It detects positive loops whenever it encounters an answer that has already been visited and which is in the SCC. In this case, the algorithm terminates returning *FALSE* to indicate the answer is unsupported. On the other hand, if the answer has been visited but is not part of the SCC, it means such answer has been produced during some other branch of query-solving and was therefore, rightfully supported and stored in the table: the algorithm terminates returning *TRUE*. Checking a non-visited answer consists of 1) marking it as visited; 2) adding it to the state of the search (stored in the *Completion Stack*); and then 3) traversing all the Delay Elements (literals) of the Delay Lists for the answer recursively checking each in turn for supportedness. Whenever an answer is determined to be unsupported, all Delay Lists containing (Delay Elements that reference) it are deleted, which may cause further simplification and iterations of ANSWER COMPLETION.

The above algorithm has been implemented within the **completion** instruction of XSB. Full performance analysis is still underway. Preliminary results indicate advantages of our heuristics: traditional benchmarks like *win/1* either do not use SIMPLIFICATION or use it seldom so that there is no overhead for ANSWER COMPLETION. A stress test that performs a large number of repetitions of Example 1 shows an overhead of at most 18%. Example 1 is actually representative of the typical situation where ANSWER COMPLETION is needed. This is so because it contains (at least) two rules for some literal (in this case  $p(X)$ ) where the first one depends on a loop through negation (rendering  $p(X)$  *undefined*) and the second one depend on a positive loop, which is unsupported. The “undefinedness” coming from the first clause is passed on to the  $p(X)$  in the body of the second one. Only ANSWER COMPLETION can then be used to clean away the delay list with  $p(X)$  from the answer coming from the second clause for  $p(X)$ . The “pathological” nature of this example follows from the, until now, XSB’s

### Algorithm Check Supported Answer(*Answer*)

```
if Answer has already been visited
  if Answer is in the SupportCheckStack return FALSE;
  else return TRUE;
else
  mark Answer as visited;
  push Answer onto the SupportCheckStack;
  mark Answer as supported_unknown;
  foreach Delay List DL for Answer
    if Answer is supported_true exit loop;
    mark DL as supported_true;
    foreach Delay Element DE in the Delay List DL
      if DL is not supported_true exit loop;
      if DE is positive and it is in the SupportCheckStack
        recursively call Check Supported Answer(Answer of DE)
        if Answer of DE is not supported_true
          mark DL as supported_false;
    if DL is supported_false
      remove DL from Answer's DLs list
      if Answer's DLs list is now empty
        delete Answer node;
        simplify away unsupported positives of Answer;
      else mark Answer as supported_true;
  if the Answer node was deleted return TRUE;
  else return FALSE;
```

Fig. 2. Algorithm CHECK SUPPORTED ANSWER.

SLG-WAM inability to rightfully detect and simplify away unsupported literals such as  $p(X)$ .

### 3. Conclusions

WFS is used in an increasing number of applications, from intelligent agents, to inheritance in object logics, to supply-chain analysis. However, the abstract complexity of WFS is a concern when embedding into the semantic core of a programming language like Prolog. Theorem 1 shows that the non-linearity of WFS can be separated from other parts of an engine for WFS; and the optimizations of the algorithm presented here, together with the preliminary performance results, underscore the suitability of WFS for general-purpose programming.

### References

1. W. Chen and D. S. Warren. Tabled Evaluation with Delaying for General Logic Programs. *JACM*, 43:20–74, 1996.
2. L. M. Pereira and G Lopes. Prospective logic agents. In *EPIA*, volume 4874 of *LNAI*, pages 73–86, 2007.
3. K. Sagonas and T. Swift. An abstract machine for tabled execution of fixed-order stratified logic programs. *TOPLAS*, 20(3):586 – 635, 1998.
4. K. Sagonas, T. Swift, and D. S. Warren. The limits of fixed-order computation. *Theoretical Computer Science*, 254(1-2):465–499, 2000.

## Appendix: Proof of Theorem 1

To help the reviewers assess the technical accuracy of Theorem 1 we include a detailed sketch of its proof, along with relevant definitions of SLG evaluation.

Theorem 1 and its proof rely on several of the definitions of SLG evaluation<sup>2</sup>. In order to make the appendix self-contained, we restate several definitions of SLG, simplified at times for ground programs and finite evaluations.

**Definition 1 (SLG Trees and Forest).** An SLG forest consists of a set of SLG trees. Nodes of SLG trees have the forms:

$$\text{Answer\_Template} :- \text{DelayList} | \text{GoalList}$$

or simply fail. In the first form, the *Answer\_Template* is an atom, *DelayList* is a set of literals and *GoalList* is a sequence of literals. The second form is called a failure node.

A node *N* is an answer when it is a leaf node for which *GoalList* is empty. If the *DelayList* of an answer is empty it is termed an unconditional answer, otherwise, it is a conditional answer.

**Definition 2 (Answer Resolution).** Let *N* be a node  $A :- D | L_1, \dots, L_n$ , where  $n > 0$ , and  $\text{Ans} = A' :- D' |$  an answer. *N* is SLG resolvable with *Ans* if  $\exists i, 1 \leq i \leq n$ , such that  $L_i$  is identical to  $A'$ . The SLG resolvent of *N* and *Ans* on  $L_i$  is:

$$(A :- D | L_1, \dots, L_{i-1}, L_{i+1}, \dots, L_n)$$

if  $D'$  is empty; otherwise the resolvent has the form:

$$(A :- D, L_i | L_1, \dots, L_{i-1}, L_{i+1}, \dots, L_n)$$

**Definition 3 (Completely Evaluated).** A set  $\mathcal{S}$  of subgoals in a forest  $\mathcal{F}$  is completely evaluated if no  $S \in \mathcal{S}$  is completed and if at least one of the following conditions holds for each  $S \in \mathcal{S}$

1. The tree for  $S$  contains an answer  $S :- |$ ; or
2. For each node  $N$  in the tree for  $S$ :
  - (a) The underlying subgoal of the selected literal of  $N$  is completed; or
  - (b) There are no applicable NEW SUBGOAL, PROGRAM CLAUSE RESOLUTION, ANSWER RETURN, NEGATION RETURN or DELAYING operations (Definition 6) for  $N$ .

**Definition 4 (Supported Answer).** Let  $\mathcal{F}$  be an SLG forest,  $S$  a subgoal in  $\mathcal{F}$ , and  $\text{Ans} = A :- \text{DelayList} |$  an answer in the tree for  $S$ . Then *Ans* is supported by  $S$  in  $\mathcal{F}$  if and only if:

1.  $S$  is not completely evaluated; or

---

<sup>2</sup> The formulation here follows T. Swift. "A New Formulation of Tabled Resolution with Delay", *Recent Advances in Artificial Intelligence*, pp. 163-177, LNAI 1695, 1999.

2. there exists an answer node  $A :- DelayList'$  of  $S$  such that every positive delay literal  $D$  is supported.

**Definition 5.** Let  $\mathcal{F}$  be an SLG forest. An atom  $S$  is successful in  $\mathcal{F}$  if some tree in  $\mathcal{F}$  has an unconditional answer  $S$ .  $S$  is failed in  $\mathcal{F}$  if  $S$  is completed and the tree for  $S$  contains no answers. A negative delay literal  $not D$  is successful (failed) in a forest  $\mathcal{F}$  if  $D$  is failed (successful) in  $\mathcal{F}$ .

**Definition 6 (SLG Operations).** Given a forest  $\mathcal{F}_n$  of an SLG evaluation of program  $P$ ,  $\mathcal{F}_{n+1}$  may be produced by one of the following operations.

1. NEW SUBGOAL: Let  $\mathcal{F}_n$  contain a non-root node

$$N = Ans :- DelayList|G, GoalList$$

where  $G$  is the selected literal  $S$  or  $not S$ . Assume  $\mathcal{F}_n$  contains no tree with root subgoal  $S$ . Then add the tree  $S :- |S$  to  $\mathcal{F}_n$ .

2. PROGRAM CLAUSE RESOLUTION: Let  $\mathcal{F}_n$  contain a root node  $N = S :- |S$  and  $C$  be a program clause  $S :- Body$ . Assume that in  $\mathcal{F}_n$ ,  $N$  does not have a child  $N_{child} = (S :- |Body)$ . Then add  $N_{child}$  as a child of  $N$ .
3. ANSWER RETURN: Let  $\mathcal{F}_n$  contain a non-root node

$$N = A :- DelayList|S, GoalList$$

whose selected literal  $S$  is positive. Let  $Ans$  be an answer node for  $S$  in  $\mathcal{F}_n$  and  $N_{child}$  be the answer resolvent of  $N$  and  $Ans$  on  $S$ . Assume that in  $\mathcal{F}_n$ ,  $N$  does not have a child  $N_{child}$ . Then add  $N_{child}$  as a child of  $N$ .

4. NEGATION RETURN: Let  $\mathcal{F}_n$  contain a leaf node

$$N = Ans :- DelayList|not S, GoalList.$$

with selected literal  $not S$ .

- (a) NEGATION SUCCESS: If  $S$  is failed in  $\mathcal{F}_n$ , then create a child for  $N$  of the form:  $Ans :- DelayList|GoalList$ .
- (b) NEGATION FAILURE: If  $S$  succeeds in  $\mathcal{F}_n$ , then create a child for  $N$  of the form  $fail$ .
5. DELAYING: Let  $\mathcal{F}_n$  contain a leaf node  $N = Ans :- DelayList|not S, GoalList$ , such that  $S$  is ground, in  $\mathcal{F}_n$ , but  $S$  is neither successful nor failed in  $\mathcal{F}_n$ . Then create a child for  $N$  of the form  $Ans :- DelayList, not S|GoalList$ .
6. SIMPLIFICATION: Let  $\mathcal{F}_n$  contain a leaf node  $N = Ans :- DelayList|$ , and let  $L \in DelayList$ 
  - (a) If  $L$  is failed in  $\mathcal{F}$  then create a child  $fail$  for  $N$ .
  - (b) If  $L$  is successful in  $\mathcal{F}$ , then create a child  $Ans :- DelayList'$  for  $N$ , where  $DelayList' = DelayList - L$ .
7. COMPLETION: Given a completely evaluated set  $\mathcal{S}$  of subgoals (Definition 3), for each  $S \in \mathcal{S}$ ,  $setMark(T, complete)$ , where  $T$  is the tree for  $S$ .
8. ANSWER COMPLETION: Given a set of unsupported answers  $\mathcal{UA}$ , create a failure node as a child for each answer  $Ans \in \mathcal{UA}$ .

**Definition 7 (SLG Evaluation).** Given a program  $P$  and goal  $G$ , an SLG evaluation  $\mathcal{E}$  is a sequence of SLG forests  $\mathcal{F}_0, \mathcal{F}_1, \dots, \mathcal{F}_\beta$ , such that:

- $\mathcal{F}_0$  is the forest containing a single tree  $G \text{ :- } | G$
- For each successor ordinal,  $n+1 \leq \beta$ ,  $\mathcal{F}_{n+1}$  is obtained from  $\mathcal{F}_n$  by an application of an SLG operation from Definition 6.

If no operation is applicable to  $\mathcal{F}$ ,  $\mathcal{F}$  is called the final forest of  $\mathcal{E}$ .

**Theorem 2.** Let  $G_{init}$  be an atomic goal to a finite ground normal program  $P$ . Under a cost model with constant time access to all subgoals, nodes, and delay literals of each SLG forest and constant time access to each clause in  $P$ , then a partial SLG evaluation  $\mathcal{E}$  that does not perform ANSWER COMPLETION can be constructed that is linear in the size of  $P$ .

(Sketch) It follows from the basic properties of SLG as shown in the literature that  $\mathcal{E}$  is finite, so that a cost function is meaningful. We consider the cost of  $\mathcal{E}$  where incremental completion is not used. The structure of the proof is

1. To indicate how the SLG operations are modified to support a particular complete scheduling strategy for  $\mathcal{E}$ ;
2. To show that the length of  $\mathcal{E}$  is linear in  $size(P)$
3. To show that each operation can be performed in constant time

*Part I: Defining the scheduling strategy.* Consider an SLG evaluation which is augmented by a queue  $Q_O$  of applicable operations, and a set  $H_F$  of non-completed subgoals that do not have a (conditional or unconditional) answer.

SLG operations are modified in the following ways to interact with  $H_F$  and  $Q_O$ .

1. When a new SLG tree for subgoal  $G_{new}$  is created, the set of PROGRAM CLAUSE RESOLUTION operations applicable to  $G_{new}$  is added to  $Q_O$ , and  $G_{new}$  is added to  $H_F$ .
2. When an SLG operation creates a new node  $N_{int}$  that is neither the root of an SLG tree nor an answer, let  $L$  be the selected literal of  $N_{int}$ , and  $G_{sel}$  the underlying atom of  $L$ .
  - (a) if  $G_{sel}$  is new to the evaluation, a NEW SUBGOAL operation is added to  $Q_O$
  - (b) if  $G_{sel}$  is not new to the evaluation,
    - i. If  $G_{sel}$  is positive, an ANSWER RETURN operation applicable for  $G_{sel}$  is added to  $Q_O$  (if an unconditional answer for  $G_{sel}$  is available in the current forest it is scheduled, otherwise an arbitrary conditional answer for  $G_{sel}$  is scheduled).
    - ii. If  $G_{sel}$  is negative
      - A. If  $G_{sel}$  is successful or failed in the current forest a NEGATION RETURN operation for  $G_{sel}$  is added to  $Q_O$
      - B. Otherwise, a DELAYING operation for  $G_{sel}$  is added to  $Q_O$
3. When an SLG operation creates a new answer  $N_{ans}$  for a subgoal  $S$ ,  $S$  is removed from  $H_F$  if it is present there, and



- (a) If  $N_{ans}$  is unconditional, and there are no other unconditional answers for  $S$  in the current forest or if  $N_{ans}$  is conditional and there are no other answers for  $S$  in the current forest, ANSWER RETURN operations are added to  $Q_O$  for each node having  $S$  as a selected positive literal
  - (b) In addition, if  $N_{ans}$  is unconditional and there are no other unconditional answers for  $S$  in the current forest, SIMPLIFICATION operations are added for each node that has  $S$  or *not*  $S$  in its delay list.
4. When a failure node is added to a node  $S$ , If  $S$  is now failed (i.e. because all of its leaf nodes are failure nodes),  $S$  is added to  $H_F$  (if it is not there), and a COMPLETION operation for  $S$  is added to  $Q_O$
  5. When a subgoal  $S$  is completed, if  $S$  is failed, SIMPLIFICATION operations for  $S$  are added to  $Q_O$ , and  $S$  is removed from  $H_F$  if it is present there.

Based on these modified operations, the initial SLG forest for  $\mathcal{E}$  consists of a tree for  $G_{init}$ , and with  $Q_O$  consisting of PROGRAM CLAUSE RESOLUTION operations applicable to  $G_{init}$ .  $\mathcal{E}$  proceeds by executing the operations in  $Q_O$  in FIFO order (actually the order does not matter). When  $Q_O$  becomes empty, a completion operation is placed on  $Q_O$  for each subgoal in  $H_F$ . If there are no such subgoals,  $\mathcal{E}$  terminates. Since each operation inserts into  $Q$  all operations that it makes applicable, the scheduling sketched here is complete.

*Part 2  $\mathcal{E}$  has a number of operations that is linear in  $size(P)$ .* We start by showing that the size of any SLG forest,  $\mathcal{F}$ , is linear in  $size(P)$ . Consider first that the number of trees in  $\mathcal{F}$  is bounded by the number of underlying subgoals of literals in  $P$ , which is linear in  $size(P)$ . Also, the total number of children of the root node of some tree in  $P$  is limited by the number of program clauses in  $P$  – again linear in  $size(P)$ . Next, let  $Node$  be node produced by PROGRAM CLAUSE RESOLUTION, with  $Lits_{Node}$  literals in its goal list. Each such literal can in principle either be resolved or delayed, leading to at most  $2^{maxlit(P)}$  children for  $N$  where  $maxlit(P)$  is the maximal number of literals in any clause in  $P$ . Accordingly,  $size(\mathcal{F})$  is linear in  $size(P)$

Next, consider that operations are scheduled only when a given node is created or a given tree is completed, each of which situations can occur only once in  $\mathcal{E}$ . We thus consider the number of operations scheduled for the creation of each type of node.

- The number of operations scheduled when creating a new tree for a subgoal  $S$  (i.e. a NEW SUBGOAL operation) is the number  $k$  of clauses with head  $S$ . Note that the total number of all NEW SUBGOAL operations is at most  $size(P)$ .
- The number of operations scheduled upon creation of an interior node with selected literal  $G_{sel}$  is 1, since it will schedule at most one NEW SUBGOAL, ANSWER RETURN, NEGATION RETURN or DELAYING operation
- The number of operations scheduled upon creation of an answer node for a subgoal  $S$  is at most  $k$  where  $k$  is the number of nodes in  $\mathcal{F}$  with  $S$  or *not*  $S$  as its selected literal or in its delay list. Note that the total cost for all such operations is linear in the size of  $P$ , since each literal on a clause of  $P$  will be resolved or simplified away at most once.
- The number of operations scheduled upon creation of a failure node is constant, since at most one COMPLETION operation will be scheduled.

- The number of operations scheduled upon completion of a node  $S$  is at most  $k$  where  $k$  is the number of nodes in  $\mathcal{F}$  with  $S$  or  $\text{not } S$  in its delay list, or  $\text{not } S$  as its selected literal. Note that the total cost of all such operations is linear in the size of  $P$ , since each literal in a clause of  $P$  will be simplified away at most once.

Since the total number of operations scheduled upon creating each type of node in the forest and the total number of operations scheduled upon subgoal completion are all linear in  $\text{size}(P)$ , the second part is shown.

*Part III: that each operation can be performed in constant time.*

Since access to  $H_F$  is assumed to be constant-time, we do not mention it in the subcases below.

- NEW SUBGOAL: Given constant-time indexing to subgoals in  $\mathcal{F}$  the NEW SUBGOAL operation, which checks whether or not a particular subgoal is in  $\mathcal{F}$  requires constant time to perform the check, and given constant-time indexing of clauses in  $P$ , the cost of scheduling each PROGRAM CLAUSE RESOLUTION operation will also require constant time.
- PROGRAM CLAUSE RESOLUTION:
  - If the operation creates an interior node  $N$ , then, given constant-time indexing to answers in  $\mathcal{F}$ , a PROGRAM CLAUSE RESOLUTION operation requires a constant time to schedule each ANSWER RETURN, NEGATION RETURN, or DELAYING operation for  $N$ .
  - Otherwise, if the operation creates an answer node for a subgoal  $S$ , constant time access to answers of  $\mathcal{F}$  allows the operation to schedule an ANSWER RETURN or NEGATION RETURN operation, if needed, for each leaf node in  $\mathcal{F}$  whose selected literal has  $S$  as its underlying atom. Given constant-time access to elements of delay lists in  $\mathcal{F}$ , each SIMPLIFICATION operation scheduled will also require constant time.
- The constant time cost of scheduling each operation by ANSWER RETURN, DELAYING, NEGATION RETURN, SIMPLIFICATION and COMPLETION operations follows from the same argument as for PROGRAM CLAUSE RESOLUTION.

Thus the application of SLG operations and their scheduling overhead requires a cost linear in  $\text{size}(P)$ . In addition, by using  $H_F$ , the scheduling of COMPLETION operations at each stage when  $Q_O$  becomes empty is proportional to the number of subgoals in  $H_F$ . Given constant-time access to subgoals in  $\mathcal{F}$  the cost scheduling these operations is linear in the size of  $H_F$ . Given a subgoal is added to  $H_F$  only once, when it is created the total cost of all scheduling of COMPLETION operations from  $H_F$ , will be linear in  $\text{size}(P)$ . Thus, the cost of all SLG operations and scheduling steps for  $\mathcal{E}$  is constant in  $\text{size}(P)$ .