Proving termination of logic programs with delay declarations

E. Marchiori and F. Teusink

Proving Termination of Logic Programs
with Delay Declarations

Elena Marchiori[1,2] and Frank Teusink[1]

[1]*CWI*

P.O. Box 94079, 1090 GB Amsterdam, The Netherlands

[2]*University of Leiden*
P.O. Box 9512, 2300 RA Leiden, The Netherlands
e-mail: {elena,frankt}@cwi.nl

Abstract

In this paper we propose a method for proving termination of logic programs with delay declarations. The method is based on the notion of *recurrent* logic program, which is used to prove programs terminating with respect to an arbitrary selection rule. Most importantly, we use the notion of bound query (as proposed by M. Bezem) in the definition of *cover*, a new notion which forms the kernel of our approach. We introduce the class of *delay recurrent* programs and prove that programs in this class terminate for all *local delay* selection rules, provided that the delay conditions imply boundedness. The corresponding method can be also used to transform a logic program into a terminating logic program with delay declarations.

1 Introduction

Delay declarations are used for the dynamic control of the selection of atoms in a derivation. The idea is that, besides the usual logic clauses, the program contains declarations of the form

<div align="center">DELAY predicate UNTIL condition</div>

Then, a selection rule is used which only selects an atom from a query, if that atom is not delayed, i.e. the condition in the delay declaration for that atom is satisfied. Delay declarations are employed in many programming systems based on logic programming, like NU-Prolog [TJ86] and Gödel [HL94]. They are important for a number of reasons: for instance, they can be used to ensure termination of the program, or to support coroutining. As a consequence, efficient algorithms can be produced from a simple logical specification augmented with suitable delay declarations. This approach reflects the idea of considering a program as consisting of two parts: logic and control.

In this paper, we study termination of logic programs with delay declarations. To illustrate how delay declarations may affect the termination behaviour of a program, consider the well-known *append*/3 program:

$$app([x|xs], ys, [x|zs]) \leftarrow app(xs, ys, zs).$$
$$app([], ys, ys).$$

and the query $app(xs, [4, 5], zs), xs = [1, 2, 3]$. This query does not terminate when the leftmost selection rule is used. However, suppose we add the following delay declaration for $append/3$:

$$\text{DELAY } app(xs, ys, zs) \text{ UNTIL } list(xs)$$

With this delay declaration, the leftmost atom in the query is delayed. Therefore, if we use a delay selection rule, only the second atom can be selected, resulting in the resolvent $app([1, 2, 3], [4, 5], zs)$. Here the atom in the query is not delayed. Moreover, this query is terminating.

The termination behaviour of a logic program with delay declarations is rather subtle. There are various aspects, sometimes unexpected, that one has to take into account. A thorough discussion of these aspects is given by Naish in [Nai92]. For instance, one would expect the delay declaration

$$\text{DELAY } app(xs, ys, zs) \text{ UNTIL } nonvar(xs) \ \lor \ nonvar(zs)$$

to ensure the termination of $append/3$. However, as illustrated by Naish, the query $app([a \,|\, T], [\,], T)$ satisfies the delay declaration, but has an infinite derivation. The fact that termination behaviour of logic programs in the context of dynamic selection rules is very subtle, is reflected also in the various methods that have been introduced, which are either based on heuristics (e.g. [LK92, MNL90]), or are rather specialized (e.g. [AL95]).

In this paper we try to tackle the problem from a different perspective. That is, we do not consider general coroutining, with all its problems, but consider the class of delay selection rules which are "local". Local selection rules are introduced in [Vie89], and correspond to selecting always in a query one of the most recently introduced atoms in the derivation from the initial query. Local selection rules behave well w.r.t. semantic information, in the following sense. If an atom of a query in a derivation is selected, then the derivation is committed to resolve that atom, and only after that atom has been completely resolved, an other atom of the query can be selected. It is this semantic property of local selection rules, which allows us to define a simple, yet powerful, method for proving termination of logic programs with delay declarations.

Our method is based on the notion of bounded query, introduced by Bezem in [Bez89, Cav89] to study termination of logic programs. We use this notion to define the central concept of our method, namely the *covers* of a body atom of a clause (query). Then, using a combination of syntactical (covers) and semantical (model) information, we define the notion of delay recurrent program. This notion is a generalization to SLD-resolution with delay selection rules, of the one of recurrent program, introduced by Bezem to study termination of logic programs w.r.t. an arbitrary selection rule. We prove that a delay recurrent program terminates for every local selection rule which selects only bounded atoms. Thus, this notion provides a method for proving termination of a logic program with delay declarations, when the delay declarations imply boundedness, i.e. if an atom satisfies its delay declaration, then that atom is bounded. Alternatively, this method can be used to find suitable delay declarations that ensure termination of goals for a given program, by choosing delay declarations which imply boundedness.

We believe that the contribution of this paper is important for at least two reasons: it provides a simple tool to reason about termination of logic programs with delay declarations, which can be also used to transform a logic program in a terminating logic program with delay declarations; moreover, it provides a new insight on the role of the selection rules when reasoning about the run-time behaviour of logic programs with delay declarations. In particular, it shows that the class of local selection rules is not only good because it supports efficient searching techniques, but also because it supports simple tools for proving termination.

The paper is organized as follows. After some preliminaries in Section 2, we present our method and the termination results in Section 3. Then, in Section 4 we give an example of proving the program $quicksort/2$ terminating in reverse order. In Section 5, we discuss some aspects of our method.

A short version of this paper appeared in [MT95].

2 Preliminaries

We shall use the following notation and terminology.

A *logic program*, called for brevity *program* and denoted by P, is a finite set of (universally quantified) clauses $H \leftarrow Q$, where Q is a query, i.e. a sequence of atoms, and H is an atom. In the following, the letters A, B indicate atoms and c a clause. For a query Q, define a Q-ground instance of a clause c to be any instance of c which grounds all the atoms of Q. Finally, c.a.s. is used as shorthand for computed answer substitution.

A sequence of atoms will also be denoted by $\tilde{A}$. As we are not interested in the order of atoms, we will sometimes treat sequences of atoms as multisets. Moreover, we will sometimes implicitly translate a sequence of atoms into a set of atoms, in order to be able to refer to elements, subsets, unions, etc. In those cases, multiplicity of atoms will be ignored, i.e. p, p will be translated into $\{p\}$. We only do this where multiplicity of atoms is not an issue.

We shall use multisets and the multiset ordering (see [Der87]). Recall that a *multiset* is an unordered collection in which the number of occurrences of each element is significant. We shall consider here the multiset ordering on multisets of natural numbers. Formally, a multiset of natural numbers is a function from the natural numbers to itself, giving the multiplicity of each natural number. Then, given the standard order $<$ on natural numbers, the ordering $<_{mul}$ on multisets is defined as the transitive closure of the replacement of a natural number with any finite number (possibly zero) of natural numbers that are smaller under $<$. Since $<$ is well-founded, the induced ordering $<_{mul}$ is also well-founded. For simplicity we shall omit in the sequel the subscript *mult* from $<_{mul}$.

A delay declaration is denoted as follows: for a predicate p of arity n a delay declaration has the form

$$\text{DELAY } p(x_1, \ldots, x_n) \text{ UNTIL } Cond(x_1, \ldots, x_n)$$

where $x_1, \ldots, x_n$ denote the arguments of p, and $Cond(x_1, \ldots, x_n)$ is a formula in some assertion language. We shall not fix the syntax of that assertion language, as it is not relevant for the sequel of the paper. The meaning of such a delay declaration is that in a query an atom $p(t_1, \ldots, t_n)$ can only be selected if the condition $Cond(t_1, \ldots, t_n)$ is satisfied. We shall assume that if an atom is selectable then all its instances are selectable too. This condition is satisfied by almost all the logic programming systems which use delay declarations. Its importance in the study of termination is crucial, and all the approaches we are aware of, for the study of properties of logic programs with delay declarations, use this assumption.

The delay declarations in a program define a class of selection rules, called *delay selection rules*. A *delay selection rule* selects an atom from a query, among those atoms which satisfy their delay declarations. If the query is non-empty and no such atom exists, no atom is selected and the query is *deadlocked*. When using delay declarations, we are only interested in SLD-derivations that are constructed using a delay selection rule. We call these derivations *delay SLD-derivations*.

3 Delay-Recurrent Programs

The aim of this paper is to define a class of programs that behave nicely with respect to termination. First, we introduce the notion of delay recurrent program. Then, we prove that, for a suitable delay declaration and a broad class of delay selection rules, every query in a delay recurrent program has only finite derivations. To this end, we use the notions of level mapping and of bounded query, introduced in [Bez89, Cav89].

Definition 3.1 (level mapping) Let P be a program. A *level mapping for P* is a function $|\,| : \mathcal{B}_P \to \mathbb{N}$ from the Herbrand base for P to the set of natural numbers. □

Thus, $|\,|$ is only defined for ground atoms. However, one can associate to a non-ground atom, the image of its set of ground instances with respect to $|\,|$:

$$\|A\| \stackrel{def}{=} \{|A'| \mid A' \text{ is a ground instance of } A\}$$

Using this, we define the notion of *bounded* atoms and queries.

Definition 3.2 (bounded query) An atom A is *bounded* (with respect to $|\ |$), if $||A||$ is finite. A query Q is bounded if all the atoms in it are bounded. $\square$

With a bounded query $Q = A_1, \ldots, A_n$ is associated the multiset $||[Q]|$ as follows:

$$||[Q]| \stackrel{def}{=} [\![max||A_1||, \ldots, max||A_n||]\!]$$

where $max||A||$ denotes the maximum of $||A||$. In the sequel, we shall often refer to $||[Q]|$ as the *level mapping* of Q.

The idea of using a level mapping to prove termination is that one proves that, in a derivation, selected atoms are always bounded and that the level mappings of the queries decrease. We can use delay declarations to ensure that only bounded atoms are selected, i.e. that the delay declarations imply boundedness.

Definition 3.3 (safe delay declaration) A delay declaration is *safe with respect to* $|\ |$ if for every atom A, if A satisfies its delay declaration then A is bounded with respect to $|\ |$. $\square$

So, by using safe delay declarations, we ensure that selected atoms are bounded. Now, we provide a method that ensures that the level mapping also decreases. For this, we use the information that selected atoms are bounded, together with the additional information provided by a model of the program. In order for an atom to be bounded, certain other atoms that originate from the same body, must have been (partially) resolved. We call these sets of atoms *covers*. To define the covers of a body atom, we need the notion of *direct covers*. Intuitively, a direct cover of an atom A in a query is a subset $\tilde{B}$ of that query, such that for some instantiation θ of the variables in $\tilde{B}$, $A\theta$ is bounded.

Definition 3.4 (direct cover) Let $|\ |$ be a level mapping. Let Q be a query, let B be an atom in Q and let $\tilde{C}$ be a subset of Q such that $B \notin \tilde{C}$. We say that $\tilde{C}$ is a *direct cover for B with respect to Q (and $|\ |$)*, if there exists a substitution θ such that $B\theta$ is bounded with respect to $|\ |$ and $Dom(\theta) \subseteq Var(\tilde{C})$.

Let H be an atom. We say that $\tilde{C}$ is a *direct cover for B with respect to $H \leftarrow Q$ (and $|\ |$)*, if there exists a substitution θ such that $B\theta$ is bounded with respect to $|\ |$ and $Dom(\theta) \subseteq Var(H, \tilde{C})$.

Finally, a direct cover $\tilde{C}$ of B is *minimal* if no proper subset of $\tilde{C}$ is a direct cover for B. $\square$

One should note that a body atom B can have zero, one, or more (minimal) direct covers. For instance when, for B to become bounded, it is necessary to instantiate a variable of B which does not occur anywhere else in the clause, B will have no direct covers. On the other hand, if B is bounded whenever H is bounded, then there exists only one minimal direct cover, namely the empty set. It is worthwhile to notice that the direct covers of an atom depend on the level mapping one chooses. For instance, consider the clause $p(x) \leftarrow p(y)$, and the two level mappings $|\ |_1$ and $|\ |_2$ such that: if s is a list then $|p(s)|_1$ is equal to its length, otherwise it is equal to 0; and $|p(s)|_2$ equal to 0 for every s. Then $p(y)$ has no direct cover w.r.t. $|\ |_1$, while it has $\emptyset$ as direct cover w.r.t. $|\ |_2$. Finally, we would like to emphasize that direct covers can be 'cyclic', in the sense that two atoms can have each other in their direct covers. Take for instance the query $p(x), q(x)$ and a level mapping $|\ |$ in which boundedness of $p(x)$ and $q(x)$ depend on x. Then, $p(x)$ will have direct cover $\{q(x)\}$ and $q(x)$ will have direct cover $\{p(x)\}$.

In the definition of cover, we take a kind of 'closure' of the direct cover relation.

Definition 3.5 (cover) Let Q be a query and let $|\ |$ be a level mapping. Let B be an atom in Q and let $\tilde{C}$ be a subset of Q. Then, $\tilde{C}$ is a *cover of B (with respect to Q and $|\ |$)*, if $\langle B, \tilde{C} \rangle$ is an element of the least set $\mathcal{C}$ ($\mathcal{C} \subseteq \mathcal{P}(Q \times \mathcal{P}(Q))$) such that

1. $\langle B, \emptyset \rangle \in \mathcal{C}$ whenever B has the emptyset as minimal direct cover, and

2. $\langle B, \tilde{C} \rangle \in \mathcal{C}$ whenever $B \notin \tilde{C}$, and $\tilde{C}$ is of the form

$$\{C_1, \ldots, C_k\} \cup \tilde{D}_1 \cup \ldots \cup \tilde{D}_k$$

such that $\{C_1, \ldots, C_k\}$ is a minimal direct cover of B in Q, and for $i \in [1..k]$, $\langle C_i, \tilde{D}_i \rangle \in \mathcal{C}$.

The notion of cover of an atom in a clause is defined analogously. $\qquad\square$

One can easily prove that the cover relation is

- 'acyclic', in the sense that if B is in a cover of A then A is not in any cover of B,

- 'monotone', in the sense that if $\tilde{C}$ is a cover of A then for all θ a subset of $\tilde{C}\theta$ is a cover of $A\theta$, and

- 'well-founded', in the sense that if there exists an atom A in Q such that A has a cover then there exists an atom B in Q such that B has an empty cover.

Using the notion of covers, we can define the class of *delay recurrent* programs.

Definition 3.6 (delay recurrent program) Let $|\ |$ be a level mapping and I an interpretation for a program P.

- A clause $c : H \leftarrow Q$ is *delay-recurrent* with respect to $|\ |$ and I if I is a model for c and for every atom A in Q, for every cover $\tilde{B}$ for A, and for every $(H, \tilde{B})$-ground instance $H' \leftarrow Q'$ of c such that H' is bounded and $I \models \tilde{B}'$, we have that

$$|[H']| > |[A']|$$

- A program P is delay-recurrent w.r.t. $|\ |$ and I if every clause is delay-recurrent with respect to $|\ |$ and I. $\qquad\square$

Knowing that a selected atom is bounded is useful, because it implies that one of the covers of that atom has been partially resolved. However, it is not enough. We need to be sure that a cover of the selected atom has been resolved completely. In order to be able to ensure this, we have to use a *local selection rule*. Local selection rules were extensively studied by Vieille in [Vie89].

Definition 3.7 (local selection rule) Let Q be a query in a derivation η, containing atoms A and B. Then A is introduced *more recently* than B, if the derivation step introducing A comes before the derivation step introducing B, in η. A is introduced *most recently*, if no atom B is introduced more recently than A.

A *local selection rule* is a selection rule that only selects *most recently* introduced atoms. $\qquad\square$

Note that, if in a query Q none of the most recently introduced atoms satisfies its delay declaration, then a local delay selection rule should deadlock on Q.

Using local selection rules, we have the following result.

Theorem 3.8 *Let P be a logic program with delay declarations. Let $|\ |$ be a level mapping and let I be an interpretation. Suppose that:*

1. *P is delay-recurrent w.r.t. $|\ |$ and I, and*

2. *the delay declarations are safe w.r.t. $|\ |$.*

Then for every query Q, every delay SLD-derivation for Q which uses a local selection rule is finite.

The proof of this theorem is rather long, and is therefore given in the Appendix. Note that we do not assume Q to be delay recurrent. We don't need to, because with the local selection rule, the atoms in Q will be resolved one at a time, without coroutining.

We conclude this section by showing that the notion of delay-recurrent program is a generalization of the notion of *recurrent programs*. This notion is due to Bezem [Bez93]. A program P is recurrent if for some level mapping $| \, |$, every ground instance $H \leftarrow A_1, \ldots, A_n$ of a clause of P satisfies the test

$$|H| > |A_i|$$

for every $i \in [1, n]$. Then we have the following result.

Lemma 3.9 *If a program P is recurrent with respect to $| \, |$ then P is delay recurrent with respect to $| \, |$ and I, for any model I of P.*

Proof: In Definition 3.6 choose I to be B_P. Consider the test $|[H']| > |[A']|$. We have that H' is bounded and Q' is (an instance of) a resolvent for H'. Thus by Lemma 2.1 of [Bez93] it follows that $|[H']| > |[Q']|$, hence $|[H']| > |[A']|$ holds.

4 An example: Quicksort

In this section, we illustrate the application of our method by means of an example. To help the reader to focus more on the method than on the example, we have chosen the well-known program *quicksort*$/2$, defined by the following set of clauses:

$$qs([x|xs], ys) \leftarrow$$
$$\qquad part(xs, x, ls, bs), \ qs(ls, sls), \ qs(bs, sbs), \ app(sls, [x|sbs], ys).$$
$$qs([], []).$$

$$part([x|xs], y, [x|ls], bs) \leftarrow x \leq y, \ part(xs, y, ls, bs).$$
$$part([x|xs], y, ls, [x|bs]) \leftarrow x > y, \ part(xs, y, ls, bs).$$
$$part([], y, [], []).$$

augmented with the clauses for *append*$/3$ given in the Introduction. Usually, the intended use of the predicate qs is that of giving it a list as first argument, in order to get a sorted permutation of that list as output in the second argument. This usage of *quicksort*$/2$ was proven to be safe (with respect to termination) e.g. in [AL95], where a proper delay declaration is chosen. Here we will show that, one can also use safely the program in its reverse, i.e. give qs a sorted list in its second argument, and it will produce all permutations of that list in its first argument. Observe that when the Prolog selection rule is used, this alternative usage of the program yields non-termination. This is the main reason why the approach of Apt and Luitjes cannot deal with this case.

We now give a level mapping for the predicates in the program, and a model. It would go too far to give a detailed account of the way we arrived at this specific level mapping. For those who are interested in techniques for finding level mappings, we refer e.g. to [DSF93]. Let $t1, \ldots, t4$ be ground terms. Then:

$|qs(t1, t2)| = tsize(t2) + 1$,
$|part(t1, t2, t3, t4) = tsize(t3) + tsize(t4)$,
$|app(t1, t2, t3)| = tsize(t3)$,
$|t1 > t2| = 0$,
$|t1 \leq t2| = 0$,

where

$$tsize(t) = \begin{cases} \text{the length of } t & \text{if } t \text{ is a list} \\ 0 & \text{otherwise} \end{cases}$$

atom	minimal direct cover	cover
$part(xs, x, ls, bs)$	$\{qs(ls, sls), qs(bs, sbs)\}$	$\left\{\begin{array}{l}qs(ls, sls), qs(bs, sbs),\\ app(sls, [x\|sbs], ys)\}\end{array}\right\}$
$qs(ls, sls)$	$\{app(sls, [x\|sbs], ys)\}$	$\{app(sls, [x\|sbs], ys)\}$
$qs(bs, sbs)$	$\{app(sls, [x\|sbs], ys)\}$	$\{app(sls, [x\|sbs], ys)\}$
$app(sls, [x\|sbs], ys)$	$\emptyset$	$\emptyset$

Figure 1: Computing covers for qs

Moreover, consider the following interpretation I:

$$
\begin{aligned}
I = \ & \{qs(t1, t2) \mid tsize(t1) = tsize(t2)\} & \cup\\
& \{part(t1, t2, t3, t4) \mid tsize(t1) = tsize(t3) + tsize(t4)\} & \cup\\
& \{app(t1, t2, t3) \mid tsize(t3) = tsize(t1) + tsize(t2)\}
\end{aligned}
$$

It is easy to check that I is a model of $quicksort/2$.

We have to prove that the clauses of $quicksort/2$ are delay recurrent with respect to this level mapping and this model. For app and $part$, this is easy to check, because they are recurrent with respect to the given level mapping. Hence the result follows from Lemma 3.9.

So, to prove the program delay recurrent, we have to check the two clauses for qs. The second clause is trivial, because it is a fact. To check the first clause, we actually have to do some work. First, we compute the minimal direct covers and covers for the atoms in the body. These are given in Figure 1. As we see, in this case every atom has a single minimal direct cover and a single cover.

Having found the covers, we can prove that the clause is delay recurrent. First of all, consider $app(sls, [x\|sbs], ys)$. A $qs(xs, ys)$-ground instance of the clause binds xs and ys to ground terms, say $t1$ and $t2$. It follows directly from the level mappings of qs and app that:

$$|[qs([x\|t1], t2)]| = tsize(t2) + 1 > tsize(t2) = |[app(sls, [x\|sbs], t2)]|$$

Secondly, $qs(ls, sls)$ has $\tilde{B} = \{app(sls, [x\|sbs], ys)\}$ as cover. A $(\tilde{B}, qs(xs, ys))$-ground instance of the clause binds xs, ys, x, sls, sbs to ground terms, say $t1, \ldots, t5$, respectively. Suppose that

$$I \models app(t4, [t3\|t5], t2).$$

Then $tsize(t2) > tsize(t4)$. But then, we have that

$$|[qs([t3\|t1], t2)]| = tsize(t2) + 1 > tsize(t4) + 1 = |[qs(ls, t4)]|$$

The proof for $qs(bs, sbs)$ is similar.

Finally, $part(xs, x, ls, bs)$ has cover $\tilde{B} = \{qs(ls, sls), qs(bs, sbs), app(sls, [x\|sbs], ys)\}$.

A $(\tilde{B}, qs(xs, ys))$-ground instance of the clause binds $xs, ys, x, sls, sbs, ls, bs$ to ground terms, say $t1, \ldots, t7$, respectively. Suppose that

$$I \models qs(t6, t4), qs(t7, t5), app(t4, [t3\|t5], t2).$$

Then $tsize(t2) > tsize(t6) + tsize(t7)$. But then we have that

$$|[qs([t3\|t1], t2)]| = tsize(t2) + 1 > tsize(t6) + tsize(t7) = |[part(t1, t3, t6, t7)]|$$

So, we have proven that $quicksort/2$ is delay recurrent with respect to $|\ |$ and I. As a result, we have that all queries will terminate, provided that a local delay selection rule is used and the delay declarations are safe. Thus, we now have to translate the boundedness information given

7

by the level mapping into delay declarations, i.e. find delay declarations for qs, $part$ and app such that if an atom is not delayed, it is bounded. For this, the following delay declarations suffice:

$$\text{DELAY } qs(xs, ys) \text{ UNTIL } list(ys)$$

$$\text{DELAY } part(xs, y, ls, bs) \text{ UNTIL } list(ls) \wedge list(bs)$$

$$\text{DELAY } app(xs, ys, zs) \text{ UNTIL } list(ys)$$

5 Discussion

In this section we discuss some aspects of our approach, and possible extensions. More precisely, we investigate the role of local selection rules in proving termination, the class of delay declarations that can be expressed using our method, and when the delay declarations do not affect the declarative semantics of the program.

5.1 Why Local Selection Rules?

In the soundness result on our method (Theorem 3.8), we restrict ourselves to local selection rules. The reason for this is that we want to use the semantic information provided by the model I. In the proof of Theorem 3.8, we use this semantic information as follows. First we observe that, when an atom A becomes selectable, some cover $\tilde{B}$ of A in the input clause that introduced A has been partially instantiated. By using the fact that a local selection rule is used, we can conclude that this cover $\tilde{B}$ has been resolved completely. As a result, we have that $I \models \forall \tilde{B}\theta$, where θ is the composition of substitutions between the node where (a generalization of) A was introduced and the node where A is selected. Finally, we use this fact to prove that the level mapping of A is strictly smaller that the level mapping of the selected atom in the resolution step that introduced (a generalization of) A.

Consider the following program:

$$\text{DELAY } p(xs, ys) \text{ UNTIL } list(xs)$$
$$p([x|xs], [x|ys]) \leftarrow p(xs, ys).$$
$$p([x|xs], [x, x|xs]) \leftarrow fail.$$
$$p([], []).$$

$$\text{DELAY } q(xs, ys) \text{ UNTIL } list(xs)$$
$$q([x|xs], [x|zs]) \leftarrow p(xs, ys), q(ys, zs).$$
$$q([], []).$$

The predicate p simply copies the list in the first argument to the second argument. However, it has a weird additional clause (clause 2) that always fails. The predicate q also copies the list in the first argument to the second argument. However, it uses p in its body, to copy xs to ys. We can prove this program delay-recurrent, using

$$|t| = \begin{cases} \text{the length of } t & \text{, if } t \text{ is a (ground) list} \\ 0 & \text{, otherwise} \end{cases}$$
$$|p(s, t)| = |s|$$
$$|q(s, t)| = |s| + 1$$
$$|fail| = 0$$

and $I = \{p(s, t) \mid |s| = |t|\} \cup \{q(s, t) \mid |s| = |t|\}$.

However, this program does not terminate for all safe selection rules. To see this, consider the following delay SLD-tree for the query $q([2, 2], vs)$.

8

$$q([2,2], vs)$$
$$[\![3]\!] \quad \{x/2, xs/[2], vs/[2|zs]\}$$
$$p([2], ys), q(ys, zs)$$
$$[\![1,2]\!] \quad \{x'/2, xs'/[], ys/[2,2]\}$$
$$fail, q([2,2], zs)$$
$$[\![0,3]\!]$$

$$\cdots$$

We see that the third query contains (a variant of) the first query. Thus, there exist infinite delay SLD-derivations for this goal obtained using a safe selection rule.

The reason for the fact that delay recurrent programs are not terminating for all delay safe selection rules is that the test for decrease of level mapping from head atom $H\theta$ to body atom $A_i\theta$ is conditional. In the case of delay recurrent programs this condition is that for some cover $\underline{B}$ for A_i, we have that $I \models \underline{B}\theta$. Translated to SLD-trees, this implies that the multisets assigned to the nodes of that tree are only guaranteed to decrease, if for every selected atom in the tree, we know that some cover of that atom is true in I. In our example, $\{p(xs, ys)\}$ is the only cover for $q(ys, zs)$ in the body of the first clause for q. In the SLD-tree, however, we select $q([2,2], zs)$ without first proving its cover, which in this case would mean first selecting $fail$. By using a local selection rule, we ensure that, before selecting an atom A_i, we first resolve a cover of that atom.

Thus, we need to restrict ourselves to the local selection rule in order to conclude that $I \models \forall \tilde{B}\theta$, which allows us to use the semantic information contained in I. This implies that our method cannot be used directly e.g. with Gödel. In fact, the Gödel selection rule selects the leftmost atoms of a query, among those which satisfy their delay declaration, even if this atom is not most recently introduced.

There is one strong argument against the use of local selection rules: they do not allow any form of coroutining. In order to prove termination with respect to selection rules that allow coroutining, we have to get rid of the restriction to local selection rules. An approach which seems quite promising, is restricting oneself to programs that do not use *speculative bindings*, a notion introduced by Naish in [Nai92]. This is something which deserves further investigation. However, we do have the impression that any method for proving termination with *full* coroutining will be either very complex, or very restrictive in its applications.

5.2 On Completeness of Delay Declarations

We have seen how delay declarations can be used to ensure termination of a logic program. One could choose strong delay declarations, like for instance DELAY $p(\tilde{x})$ UNTIL *false*, which certainly imply termination. However, the resulting program would not be very interesting, since it yields no c.a.s.'s. To ensure that the delay declaration is not too strong, one has to guarantee that the declarative semantics of the program is preserved. This is specified in the following definition.

Definition 5.1 (complete delay declaration) Let P be a program and let I be the least Herbrand model for P. Let $\mathcal{D}$ be a set of delay declarations for P. We say that $\mathcal{D}$ is *complete* w.r.t. P if every atom in I has a successful delay SLD-derivation in $P \cup \mathcal{D}$.

A sufficient condition for completeness of a delay declaration w.r.t. P is that every ground atom which is in I is deadlock free. An atom is deadlock free if all its finite derivations do not end in a non-empty query which contains only atoms that do not satisfy their delay declarations. Then, the following result holds.

9

Lemma 5.2 *Let P be a program and let I be the least Herbrand model for P. Let $\mathcal{D}$ be a set of delay declarations for P. Suppose that every atom A of I is deadlock-free. Then $\mathcal{D}$ is complete with respect to P.*

Recently, the topic of deadlock-freedom of programs with delay declarations has been studied in [AL95, CD94, EG96, MT96]. The methods there introduced can be applied to prove that every atom of I is deadlock-free.

5.3   On Expressiveness of Delay Declarations

In Gödel, one can use the predicate *nonvar* in delay declarations. For instance, the following delay declaration is used for the predicate *app* defined by the program given in the Introduction:

$$\text{DELAY } app(xs, ys, zs) \text{ UNTIL } nonvar(xs) \ \lor \ nonvar(zs)$$

When this delay declaration is used, an atom $app(s, t, u)$ is not selected until either s or u is a non-variable term.

We cannot deal with these kinds of delay declarations. The reason is that in our definition of delay recurrent programs, the notion of level mapping we use is the one used in the definition of recurrent programs. In this definition, the level mapping $|A|$ for ground atoms A is defined by a (total) function from $\mathcal{B}_P$ to $\mathbb{N}$, whereas the level mapping $||B||$ for non-ground atoms B is defined as the maximum of the level mappings of all its ground instances. Thus $||\ ||$ is a partial function, because the set of level mappings of ground instances can be unbounded. As a consequence, when taking the level mapping of an atom $p(l)$ to be the length of list l, the atom $p([x|xs])$ contains a non-variable term, but $||p([x|xs])||$ is undefined because xs can be instantiated with an arbitrary large ground list. Thus, an atom $app([x|xs], ys, zs)$ is not bounded, while it satisfies the condition of the delay declaration. Terms which behave well with respect to a level mapping have been studied for instance in [BCF94], where they are called rigid.

As the *append*/3 example given in the Introduction shows, the termination behaviour of "delay until nonvar" is poorly understood. As far as we can see now, a method handling the *nonvar* delay predicate would also be significantly more complex (or, alternatively, weaker), than our method. All in all, the problems with the *nonvar* delay predicate were enough for us to decide not to deal with it at this point. As a final remark we would like to note that, if one browses through the Gödel manual, it seems that our method is severely handicapped by not being able to handle *nonvar*, because most delay declarations in example programs use *nonvar*. One should note however, that these programs are not guaranteed to terminate for all goals (not even when the leftmost undelayed selection rule is used). To be fair, the Gödel manual only states that the delay declarations can be used to *assist* termination. On the other hand, our method guarantees termination, be it that the delay declarations will be more restrictive.

5.4   On programs with negation

It seems that our method can be easily extended to deal with logic programs with negation. We sketch briefly how this could be done. One can extend the procedure for resolving negated atoms to the case of delay selection rules, simply considering a form of (abnormal) termination, which arises when a tree for $\neg A$ is finite but contains at least one leaf consisting of delayed literals. In such a case $\neg A$ has no resolvent; it ends in deadlock. Then the definition of level mapping can be extended to negated atoms, simply by defining $|\neg A| = |A|$. Finally, in Definition 3.6 of delay-recurrent program, the model I should be replaced by some model containing suitable semantic information.

6   Related Work

Let us now relate our approach to other work on termination with respect to dynamic selection rules.

The paper which helped us to understand the problems in reasoning about the termination of logic programs with delay declarations, is [Nai92]. In this paper, L. Naish investigates how termination of a conjunction of queries can be established, under the hypothesis that the execution of each query does terminate. However, he does not propose ready to use methods for proving programs terminating. In his paper, Naish argues that the use of modes is crucial to reasoning about termination. To support this claim, he gives a number of useful observations on the termination behaviour of a program with delay declarations, which emphasize how subtle is this behaviour, and how difficult it is to prove termination, when dealing with general coroutining. Towards the end of the paper, Naish suggests that the existence of 'speculative bindings' are an important complicating factor when reasoning about termination. It might be the case that in absence of these speculative bindings, we can generalize our method to non-local delay selection rules.

Another recent contribution to the subject of termination with respect to delay declarations, is [LK92]. In this paper, S. Lüttringhaus-Kappel discusses a non-deterministic scheme for finding delay declarations that ensure termination. First, he presents an algebra of 'when' declarations. This algebra is more expressive than the class of delay declarations we can handle, basically because we cannot handle *nonvar* predicates. The scheme itself is very general; it is meant as a basis for practical implementations, using heuristics and partial evaluation to replace non-deterministic choices. The results of an existing implementation look quite promising. On the other hand, as the scheme is very general, it does not give much insight in the problem of termination itself. Another problem is that one has to prove that a program is 'safe' (not the notion used in this paper), which is quite difficult, the more because there are no methods for doing this.

A very recent paper by K.R. Apt and I. Luitjes [AL95], stimulated us to work on our approach. In this paper, they discuss verification of logic programs with respect to dynamic selection rules. In one section they discuss the problem of termination. The approach they take is more general than ours, in the sense that they do not restrict to local selection rules. As a consequence, they need to impose strong restrictions on the class of programs they consider. One restriction in this work is that the termination results are stated in terms of termination with respect to LD-resolution. Thus, it can only discuss termination with respect to dynamic selection rule of programs which are known to terminate with respect to leftmost selection rule.

It is clear that most of the programs we can prove terminating with our method, can also be proven to be terminating by a static reordering of bodies of program clauses. We think however that the use of covers has a number of advantages. First of all, with covers we have a systematic approach for finding static orderings that ensure termination, which is more efficient than simply checking all permutations of body atoms. Secondly, our method does not impose an order on body atoms. If one fixes the order of body atoms in order to ensure termination, one looses the freedom to let a compiler or optimizer fix some order. Instead, the covers computed in our method form a concise representation of *all* orderings of body atoms that ensure termination. This information can be fed to a compiler or optimizer, as a constraint on the orderings of bodies it may choose. Finally, there exist programs that can be proven terminating with our approach, which are not (easily) proven terminating with a static approach.

7 Conclusion

In this paper we introduced a simple method for proving termination of logic programs with delay declarations. The method is based on the new notion of cover, which is used to describe the inter-relation among the atoms of a clause that can be caused by the dynamic scheduling. Covers are used to define the class of delay recurrent programs. We proved that all derivations of a delay-recurrent program are finite, when the selection rule is local delay, i.e. it selects at each resolution step one atom which satisfies its delay declaration, among those atoms most recently introduced. We discussed advantages and limitations of this last condition on the selection rule. We intend to continue investigating other conditions under which we can relax the restriction to local selection rules, although we think that such methods are necessarily either much more complex or applicable to much smaller classes of programs.

Acknowledgements

References

[AL95]   K. R. Apt and I. Luitjes. Verification of logic programs with delay declarations. In *Proceedings of the Fourth International Conference on Algebraic Methodology and Software Technology, (AMAST'95)*, Lecture Notes in Computer Science, Berlin, 1995. Springer-Verlag. Invited Lecture. In preparation.

[BCF94]   A. Bossi, N. Cocco, and M. Fabris. Norms on terms and their use in proving universal termination of logic programs. *Theoretical Computer Science*, 124:297–328, 1994.

[Bez89]   M. Bezem. Characterizing termination of logic programs with level mappings. In E. L. Lusk and R. A. Overbeek, editors, *Proceedings of the North American Conference on Logic Programming*, pages 69–80. The MIT Press, 1989.

[Bez93]   M. Bezem. Strong termination of logic programs. *Journal of Logic Programming*, 1 & 2(15):79–97, 1993.

[Cav89]   L. Cavedon. Continuity, consistency, and completeness properties for logic programs. In G. Levi and M. Martelli, editors, *Proceedings of the International Conference on Logic Programming*, pages 571–584. The MIT Press, 1989.

[CD94]   P. Chambre and P. Deransart. Towards a proof method of non-suspension of Concurrent Constraint Logic Programs. In F. de Boer and M. Gabbrielli, editors, *Proceedings of the W2 Post-Conference Workshop, Int. Conf. on Logic Programming*, Amsterdam, 1994. Free Univerity.

[Der87]   N. Dershowitz. Termination of rewriting. *Journal of Symbolic Computation*, 8:69–116, 1987.

[DSF93]   S. Decorte, D. De Schreye, and M. Fabris. Automatic inference of norms: a missing link in automatic termination analysis. In *Proceedings of the International Logic Programming Symposium*, pages 420–436, 1993.

[EG96]   S. Etalle and M. Gabbrielli. Layered Modes. In F. de Boer and M. Gabbrielli, editors, *Proc. of the JICSLP post-conference workshop on Verification and Analysis of Logic Programs*. University of Pisa, 1996.

[HL94]   P.M. Hill and J.W. Lloyd. *The Gödel Programming Language*. The MIT press edition, 1994.

[LK92]   Stefan Lüttringhaus-Kappel. *Laziness in Logic Programming*. PhD thesis, Mathematisch-Naturwissenschaftlichen Fakultät der Rheinischen Friedrich-Wilhelms Universität Bonn, Bonn, 1992.

[MNL90]   K. Marriott, L. Naish, and J.L. Lassez. Most specific logic programs. *Annals of mathematics and artificial intelligence*, 1(2), 1990.

[MT95]   E. Marchiori and F. Teusink. Proving termination of logic programs with delay declarations. In J. Lloyd, editor, *Proceedings of the 1995 International Symposium on Logic Programming*, pages 447–464. The MIT Press, 1995.

[MT96]   E. Marchiori and F. Teusink. On deadlock freedom of logic programs with dynamic schedulin. In F. de Boer and M. Gabbrielli, editors, *Proc. of the JICSLP post-conference workshop on Verification and Analysis of Logic Programs*. University of Pisa, 1996.

[Nai92]   L. Naish. Coroutining and the construction of terminating logic programs. Technical Report 92/5, Department of Computer Science, University of Melbourne, 1992.

[TJ86]   J. Thom and J.Zobel. NU-prolog reference manual. Technical Report 86/10, University of Melbourne, 1986.

[Vie89]   L. Vieille. Recursive query processing: the power of logic. *Theoretical Computer Science*, 69(1):1–53, 1989.

A   Proof of Theorem 3.6

In this appendix, we give a proof of Theorem 3.8. This proof consists of a number of preliminary results, which yields Theorem 3.8 in the form of Corollary A.11.

Bezem proved that for a recurrent program, all bounded queries will terminate for all selection rules. In essence, the termination proof for recurrent programs proves that, in every derivation for a bounded goal, the level mapping of a goal is greater than the level mapping of its resolvent. As the order on multisets is well-founded, it follows that the derivation has to be finite. The termination proof for delay-recurrent programs is similar to the one for recurrent programs, except that we use the additional information provided by a model of the program. Moreover, we use the fact that the selected atom is always bounded.

Recall that we are working under following two assumptions:

- the delay declarations are safe and

- the selection rules are local.

First, we state a useful persistence property of being delay-recurrent.

Lemma A.1 *A head-instance of a delay-recurrent clause is delay-recurrent.*

Note that the above result does not hold when a generic instance is considered. For instance, it could happen that an unbounded atom occurring in the body of the clause becomes bounded by instantiating some variables occurring only in that atom.

Next, we fix some terminology in order to be able to speak about the queries and atoms in a delay SLD-derivation. Suppose we have a delay SLD-derivation containing a query $Q = A, N_1, \ldots, N_k$ and its resolvent $Q' = B_1, \ldots, B_l, N_1, \ldots, N_k\theta$. We say that, for $i \in [1..k]$, the atom N_i in Q is the *parent* of the atom $N_i\theta$ in Q' (or vice versa, $N_i\theta$ is a *child* of N_i). Analogously, we call Q the parent of Q' (and Q' the child of Q). Moreover, A is the *direct generator* of the atoms $B_1, \ldots, B_l$ in Q'. The *predecessor (resp. descendent)* relation (on atoms) is the transitive closure of the parent (resp. child) relation. We call with the same name the analogous relations defined on queries instead of atoms. Finally, an atom A (in a node Q) is the *generator* of an atom A' (in a node Q'), if

- A is the direct generator of A', or

- $A = A'$ and $Q = Q'$ is the first node of the derivation, or

- A is the direct generator of an atom A'' in the resolvent of Q, and A'' is a predecessor of A'.

Given an atom A, we use $gen(A)$ to denote the generator of A.

Before we even begin with the termination proof, let us first state some properties of local selection rules which will prove to be essential in the termination proof.

Lemma A.2 *Let η be a derivation using a local selection rule. Let Q and Q' be queries in η such that Q is a predecessor of Q', and let A and A' be their respective selected atoms. Let Q_g and Q'_g be the queries that contain $gen(A)$ and $gen(A')$, respectively. If Q'_g is a predecessor of Q, then Q'_g is a predecessor of Q_g.*

Proof: Because the local selection rule is used, the selected atoms are always those that "have been introduced in the derivation most recently". That is equivalent to saying that their generators are closest. $\square$

Thus, atoms are selected according to a LIFO (last in, first out) scheme.

Definition A.3 Let η be a derivation using a local selection rule, with queries $Q_0, Q_1, \ldots$, input clauses $H_0 \leftarrow \tilde{A}_0, H_1 \leftarrow \tilde{A}_1, \ldots$ and mgu's $\theta_0, \theta_1, \ldots$. A *segment* η_j^i of η consists of queries $Q_i, \ldots, Q_j$, input clauses $H_i \leftarrow \tilde{A}_i, H_{j-1} \leftarrow \tilde{A}_{j-1}$ and mgu's $\theta_i, \ldots, \theta_{j-1}$. We define $\theta_j^i = \theta_i \ldots \theta_{j-1}$. We call a segment η_j^i *complete* if

1. $i < j$,

2. for all queries Q_k with $k \geq j$, for all atoms A' in Q_k, $gen(A')$ is not contained in a query in $Q_i, \ldots, Q_{j-1}$, and

3. for all queries Q_k with $i < k < j$, the generator of the selected atom in Q_k is contained in a query in $Q_i, \ldots, Q_j$. $\qquad\square$

Lemma A.4 *Let η be a derivation using a local selection rule. Let η_j^i be a complete segment of η. Let A be the selected atom in Q_i. Then $\theta_j^i|_{Var(A)}$ is a c.a.s. for A.*

Proof: We prove the claim by induction on the length $l = (j - i) + 1$ of η_j^i. Note, that the length of a complete segment is at least 2. For $l = 2$, we have that $\theta_j^i = \theta_i$. Consider the input clause $H_i \leftarrow \tilde{A}_i$. Because η_j^i is a complete segment, $\tilde{A}_{j-1}$ must be the empty query. But then, from $i = j - 1$ it follows that $\tilde{A}_i$ is the empty query. Because $A\theta_i = H\theta_i$, we have that A has the empty query as resolvent, using $H_i \leftarrow \tilde{A}_i$ as input clause, and θ_i as mgu. It follows that θ_i is a c.a.s. for A. Now, assume that the claim holds for all segments of length less than l.

We have to prove the claim for segments of length l. Consider the input clause $H_i \leftarrow \tilde{A}_i$ for Q_i. Let $\tilde{A}_i = A_1, \ldots, A_m$. Because η uses a local selection rule, η_j^{i+1} can be divided into m complete segments, one for every A_k, for $k \in [1..m]$. Without loss of generality, let us assume that, for $k \in [1..m]$, the complete segment for A_k is the k-th complete segment in η_j^{i+1}. Let, for $k \in [1..m]$, μ_k be the composition of the mgu's of the k-th complete segment. By induction hypothesis, $\mu_k|_{Var(A_k)}$ is a c.a.s. for A_k. But then, $(\mu_1, \ldots, \mu_m)|_{Var(A_1, \ldots, A_m)}$ is a c.a.s. for the query $A_1, \ldots, A_m$. Because we have that $A\theta_i = H_i\theta_i$, it follows that $\theta_j^i|_{Var(A)} = (\theta_i\mu_1 \ldots \mu_m)|_{Var(A)}$ is a computed answer substitution for A. $\qquad\square$

In the sequel, we assume that $\theta_j^i = \epsilon$ if $i \geq j$.

Lemma A.5 *Let η be a derivation using a local selection rule, with queries $Q_0, Q_1, \ldots$, input clauses $H_0 \leftarrow \tilde{A}_0, H_1 \leftarrow \tilde{A}_1, \ldots$ and mgu's $\theta_0, \theta_1, \ldots$. Let η_j^i be a segment of η such that the generator of the selected atom in Q_j is the selected atom in Q_i. Let $\mu = \theta_j^{i+1}|_{Var(\tilde{A}_i\theta_i)}$. Then, μ is a c.a.s. for a subquery of $\tilde{A}_i\theta_i$.*

Proof: We prove the claim by induction on the length $l = (j - i) + 1$ of η_j^i. For $l = 1$, we have that $\theta_j^{i+1} = \epsilon$, and therefore the claim holds, if we take the subquery to be the empty query. Now, consider $l > 1$ and assume that the claim holds for all segments smaller than l.

First of all, a prefix of η_j^i can be handled directly using induction. Suppose that for some Q in $Q_{i+1}, \ldots, Q_{j-1}$, the generator of the selected atom in Q is Q_i. Then, let k be the greatest number in $[i + 1..j - 1]$ such that the generator of the selected atom in Q_k is Q_i. Otherwise, let $k = i$. Now, let $\mu' = (\theta_{i+1} \ldots \theta_{k-1})\_Var(\tilde{A}_i\theta_i))$. By induction hypothesis, μ' is a c.a.s. of a subquery of $\tilde{A}_i\theta_i$.

We now deal with the remainder of η_j^i, i.e. θ_j^k. As the generator of the selected atom of Q_k is in Q_i, we know that the selected atom in Q_k is of the form $A\mu'$, where A is an element of $\tilde{A}_i\theta_i$. Moreover, the generator of the selected atom in Q_j is also in Q_i. But then, from Lemma A.2, it follows that η_j^k is a complete segment and therefore, by Lemma A.4, $\theta_j^k|_{Var(A\mu')}$ is a c.a.s. for $A\mu'$. Let $\mu'' = \theta_j^k|_{Var(A\mu')}$

Composing μ' and μ'' we have that $(\mu'\mu'')|_{Var(\tilde{A}_i\theta_i)} = \mu'\mu''$ is a c.a.s. for a subquery of $\tilde{A}_i\theta_i$. $\square$

Having stated some results on local selection rules, we are now able to prove termination of delay recurrent programs.

Theorem A.6 *Let P be a program that is delay-recurrent w.r.t. $|\ |$ and I. Let τ be a delay SLD-tree for P (w.r.t. $|\ |$) having an atom as root. Then τ is finite.*

15

The termination proof for delay recurrent programs will proceed as follows. First, we label every query Q in the SLD-tree with a multiset $||[Q]||_d$. Then, we show that for every query Q in the tree, and every resolvent R of Q, $||[R]||_d$ is smaller than $||[Q]||_d$ in the multiset ordering.

So, let us define the procedure to label the nodes of an SLD-tree.

Definition A.7 ($||[Q]||_d$) Let $|\ |$ be a level mapping for P. Let τ be a delay SLD-tree for P having a bounded atom as root. Let A be an atom in some query in τ. Then,

$$||A||_d = \begin{cases} ||A|| & \text{, if A} = gen(\text{A}) \\ ||gen(A)|| - 1 & \text{, otherwise} \end{cases}$$

Let $Q = A_1, \ldots, A_n$ be a node in τ. Then,

$$||[Q]||_d = [\![max||A_1||_d, \ldots, max||A_n||_d]\!]$$

$\square$

The idea behind the label $||[Q]||_d$ is that it is, in some sense, a 'safe' estimation of $||[Q]||$. We have to use such an estimation because the atoms in an SLD tree are not necessarily bounded. Therefore $||Q||$ is not defined for all queries in the SLD tree. To solve this problem, we define $||A||_d$ in terms of $||B||$, where B is the generator of A. In the case of delay SLD trees, $||B||$ is defined, because in an delay SLD tree the selected atom of a query is guaranteed to be bounded. In the following lemma, we prove that $||[A]||_d$ is 'safe', in the sense that, for all selected atoms in the queries of an delay SLD tree, $||A||_d \geq ||A||$.

Lemma A.8 *Let P be a program that is delay-recurrent w.r.t. $|\ |$ and I. Let τ be a delay SLD-tree for P (w.r.t. $|\ |$) having an atom as root. Let Q be a query of τ and let A be its selected atom. Then, $||A||_d \geq ||A||$.*

Proof: If the generator of A is A itself, then the claim follows immediately from Definition A.7.

Otherwise, let η be the branch of τ that contains Q, and let Q be the j-th query in η (i.e. $Q = Q_j$). Let B be the generator of A, and let Q_i be the node that contains B. Now, consider the segment η_j^i of η, and let $\mu = \theta_j^i|_{Var(\tilde{A}_i\theta_i)}$.

By Lemma A.5, we have that μ is a c.a.s. for a subquery, say $\tilde{B}$, of $\tilde{A}_i\theta_i$. This implies that $I \models \forall(\tilde{B}\mu)$. Let A' be the atom of $\tilde{A}_i\theta_i$ such that A is its descendant. Since A is bounded, a subquery, say $\tilde{C}$, of $\tilde{B}$ is a cover for A'. Let $\beta = \mu|_{Var(\tilde{C})}$. Then we have that $I \models \tilde{C}\beta\rho$, for every ρ which grounds $\tilde{C}\beta$.

By the hypothesis that P is delay recurrent and by Lemma A.1 it follows that $H\theta_i \leftarrow \tilde{A}_i\theta_i$ is delay-recurrent. Then:

$$||[H\theta_i\beta\rho]|| > ||[A'\beta\rho]||. \tag{1}$$

But then, because B is bounded and $H\theta_i = B\theta_i$, it follows that

$$||[B]|| \geq ||[B\theta_i]|| = ||[H\theta_i]|| \geq ||[H\theta_i\beta\rho]|| > ||[A'\beta\rho]||$$

for every ρ. From $A'\beta$ bounded it follows that $||[B]|| > ||[A'\beta]||$. Finally, the claim follows from the fact that $A = A'\beta\gamma$ for some γ and $||[B]|| = ||[A]||_d$. $\square$

Having proven that $||[Q]||_d$ is a good approximation for $||[Q]||$, we now prove that the label of a query is greater than the label of its resolvent. To this end, we use the following result by Bezem.

Proposition A.9 *Let Q be a bounded query and let θ be a substitution. Then,*

1. *$Q\theta$ is bounded.*

2. *$||[Q]|| \geq ||[Q\theta]||$*

Lemma A.10 *Let P be a program that is delay-recurrent w.r.t. $||$ and I. Let τ be a delay SLD-tree for P (w.r.t. $||$) having an atom as root. Let Q be a query of τ and let R be a direct descendant of Q. Then $||[R]||_d$ is smaller than $||[Q]||_d$ in the multiset ordering.*

Proof: Suppose for simplicity that the leftmost atom of Q is selected. Then, $Q = A, \tilde{B}$ and $R = (\tilde{C}, \tilde{B})\theta$, with $H \leftarrow \tilde{C}$ input clause and $\theta = mgu(H, A)$.

For every B in $\tilde{B}$ we have that $||[B\theta]||_d = ||[B]||_d$. Moreover, for every C in $\tilde{C}\theta$, we have from Definition A.7 that $||[C]||_d = ||[H\theta]|| - 1$. Thus, to prove the claim, it is sufficient to show that $||[A]||_d > ||[H\theta]|| - 1$. We have that:

$$
\begin{array}{rll}
||[A]||_d & \geq & ||[A]|| \qquad \text{from Lemma A.8} \\
& \geq & ||[A\theta]|| \qquad \text{from Proposition A.9} \\
& = & ||[H\theta]|| \qquad \text{from } \theta = mgu(H, A) \\
& > & ||[H\theta]|| - 1
\end{array}
$$

$\square$

Proof: (of Theorem A.6)
From Lemma A.10 (recall that we work with safe selection rules) and the fact that the multiset order of Definition A.7 is well-founded. $\square$

Then we obtain Theorem 3.8 in the form of the following corollary.

Corollary A.11 *Let P be a program that is delay-recurrent w.r.t. $||$ and I. Let τ be a delay SLD-tree for P (w.r.t. $||$). Then τ is finite.*

Proof: Let $Q = A_1, \ldots, A_n$ be the root of τ. We prove the result by induction on n. If $n = 1$ then the result follows by Theorem A.6. Suppose now $n > 0$. If every atom of Q is not bounded then the result follows immediate. Otherwise let A_i be the selected atom. Then from Theorem A.6 we have that the delay SLD-tree for A_i is finite. Let $\theta_1, \ldots, \theta_m$ be its computed answers. Then from the assumption that the selection rule is local, it follows that τ can be split into $m + 1$ delay SLD-trees, τ_1 for A_i and the others τ_k for $(Q - \{A_i\})\theta_k$, for $k \in [1, m]$. Then the result follows by the application of the induction hypothesis to the τ_k's. $\square$