

(Co)recursion in Logic Programming: Lazy vs Eager*

JÓNATHAN HERAS

School of Computing, University of Dundee, UK
(e-mail: jonathanheras@computing.dundee.ac.uk)

EKATERINA KOMENDANTSKAYA

School of Computing, University of Dundee, UK
(e-mail: katya@computing.dundee.ac.uk)

MARTIN SCHMIDT

Institute of Cognitive Science, University of Osnabrück, Germany
(e-mail: martisch@uos.de)

submitted 1 January 2003; revised 1 January 2003; accepted 1 January 2003

Abstract

CoAlgebraic Logic Programming (CoALP) is a dialect of Logic Programming designed to bring a more precise compile-time and run-time analysis of termination and productivity for recursive and corecursive functions in Logic Programming. Its second goal is to introduce guarded lazy (co)recursion akin to functional theorem provers into logic programming. In this paper, we explain lazy features of CoALP, and compare them with the loop-analysis and eager execution in Coinductive Logic Programming (CoLP). We conclude by outlining the future directions in developing the guarded (co)recursion in logic programming.

KEYWORDS: Logic Programming, Recursion, Corecursion, Termination, Productivity, Guardedness.

1 Introduction

Logic Programming (LP) was conceived as a *recursive* programming language for first-order logic. Prolog and various other implementations of LP feature *eager* derivations, and therefore *termination* has been central for logic programming (de Schreye and Decorte 1994). However, unlike e.g. functional languages, LP has not developed an operational semantics supporting explicit analysis of termination. In typed programming languages like Coq or Agda, it is possible to introduce syntactic (static) checks that ensure *structural* recursion, and hence termination of programs at run-time. In Prolog, there is no support of this kind.

Example 1.1 (BitList) *Consider the following recursive program that defines lists of bits.*

```
1. bit(0) ←  
2. bit(1) ←  
3. bitlist([]) ←  
4. bitlist([X|Y]) ← bit(X), bitlist(Y)
```

* The work of the first two authors was supported by EPSRC Grant EP/K031864/1.

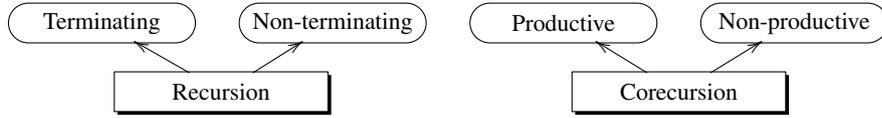


Fig. 1. Distinguishing well-founded and non-well-founded cases of recursion and corecursion.

It is a terminating program, however, if the order of clauses (3) and (4), or the order of atoms in clause (4) is accidentally swapped, the program would run into an infinite loop.

This example illustrates that non-terminating (co)recursion is distinguished only empirically at run-time in LP. This distinction is not always accurate, and may depend on searching strategies of the compiler, rather than semantic meaning of the program.

Coinductive Logic Programming (CoLP) (Gupta et al. 2007; Simon et al. 2007) has been introduced as a means of supporting *corecursion* in LP. A representative example of coinductive programming is to reason about an infinite data structure, for example an infinite stream of bits.

Example 1.2 (BitStream) Given the definition of bits as in Example 1.1, an infinite stream of bits is defined as:

$$1. \text{stream}([X|Y]) \leftarrow \text{bit}(X), \text{stream}(Y)$$

Note that unlike **BitList**, we no longer have the base case for recursion on *stream*.

The tradition (Coquand 1994) has a dual notion to termination for well-behaving *corecursion* – and that is *productivity*. If termination imposes the condition that any call to an inductively defined predicate like `bit` must terminate, then productivity requires that every call to a coinductive predicate like `stream` must *produce* some partially *observed* structure in a finite number of steps. E.g. calling `stream(X)?`, the program must compute an answer `[0|Y]` observing the component `0` in finite time. Moreover, the productivity imposes a second condition: the computation must be able to proceed corecursively, e.g. in our example, the condition is for `Y` to be an infinite productive datastructure. This situation is explained in e.g. (Abel et al. 2013; Bertot and Komendantskaya 2008).

CoLP deals with programs like **BitStream** by using a combination of eager evaluation, SLD-resolution and loop analysis. In simplified terms, for a goal `stream(X)?` the resolvent loop detection would allow to return an answer `X=[0|X]`; by observing the “regular” pattern in resolvents involving Clause (1) in the derivations. Similarly to standard (recursive) LP, non-terminating cases of corecursion (where no regular loop can be found) are not formally analysed in CoLP.

Example 1.3 (BadStream) *BadStream* is not productive; that is, it would be executed infinitely without actually constructing a stream.

$$1. \text{badstream}([X|Y]) \leftarrow \text{badstream}([X|Y])$$

A different case of corecursion is the below example, which is productive, but cannot be handled by CoLP loop detector, as the stream it defines is not regular.

Example 1.4 (TakeFirstN) The program *TakeFirstN* defines the stream of natural numbers,

and allows to construct a list with the first n elements of the stream by calling *taken*.

1. $from(X, [X/Y]) \leftarrow from(s(X), Y)$
2. $take(0, Y, []) \leftarrow$
3. $take(s(X), [Y/Z], [Y/R]) \leftarrow take(X, Z, R)$
4. $taken(N, X) \leftarrow from(0, Y), take(N, Y, X)$

In CoLP, calls to e.g. $taken(s(s(0)), X)$? fall into infinite computations that are not handled by the loop detection procedure. Similar to how LP would be unable to handle **BitList** with swapped atoms in clause (4) though in principle the program describes a well-founded inductive structure, CoLP would not be able to handle **TakeFirstN** although it is a perfectly productive stream. For the query *taken*, it is intuitively clear that, the construction of the first n elements of the stream should take a finite number of derivation steps.

Coalgebraic Logic Programming (CoALP) (Komendantskaya and Power 2011; Komendantskaya et al. 2014a) gives a new (coalgebraic) operational semantics for LP; and in particular it offers new methods to analyse termination and productivity of logic programs. Using CoALP, we present here a coherent operational treatment of recursion and corecursion in LP, and discuss new methods to distinguish well-founded and non-well-founded cases of (co)recursion in LP, as outlined in Figure 1. Unlike Prolog or CoLP, CoALP is a first lazy dialect of logic programming; and it features guarded (co)recursion akin to structural recursion and guarded corecursion in e.g. Coq or Agda (Coquand 1994; Abel et al. 2013). The current implementation of CoALP in parallel language Go is available in (Komendantskaya et al. 2014b); and is tested on a few benchmarks in this paper. Here, we abstract from some of the technical details available in (Komendantskaya et al. 2014a) and from implementation details available in (Komendantskaya et al. 2014c) and give a higher-level discussion of the issues of termination and productivity in LP.

The rest of the paper is structured as follows. In Section 2, we explain the role of laziness in semantics and implementation of CoALP; in Section 3, we discuss the effect of guarded corecursion. Section 4 is devoted to discussion of our current work on soundness properties for corecursive logic programming.

2 Lazy Corecursion in Logic Programming

CoALP uses the standard syntax of Horn-clause logic programming (Lloyd 1987), but offers a new derivation algorithm in place of the SLS-resolution. One of the main distinguishing features of CoALP is its laziness. To our knowledge, it is the first lazy dialect of logic programming. The issue is best explained using the following example:

Example 2.1 Given the program **BitList** and the query $bitlist([X/Y])$, the standard algorithm of SLD-resolution (Lloyd 1987) will eagerly attempt to find a derivation, e.g.:

$$bitlist([X/Y]) \longrightarrow bit(X), bitlist(Y) \xrightarrow{X=0} bitlist(Y) \xrightarrow{Y=[]} \square$$

For the program **BitStream** this will give rise to an infinite SLD-derivation:

$$stream([X/Y]) \longrightarrow bit(X), stream(Y) \xrightarrow{X=0} stream(Y) \xrightarrow{Y=[X1/Y1]} stream([X1/Y1]) \dots$$

In the above setting, there is no natural place for laziness, as ultimately the strong side of the procedure is a fully automated proof search. Fibrational coalgebraic operational semantics of LP

induced by A , and let C be a clause $H \leftarrow B_1, \dots, B_n$. Then, the goal G' is coinductively derived from G and C using the mgu θ if the following conditions hold:

- ★ $Q(\bar{t})$ is a node in T .
- ★★ θ is an mgu of $Q(\bar{t})$ and H .
- ★★★ G' is given by the (coinductive) tree $T\theta$ with the root $A\theta$.

Definition 2.3 A coinductive derivation of $P \cup \{G\}$ consists of a sequence of goals $G = G_0, G_1, \dots$ and a sequence $\theta_1, \theta_2, \dots$ of mgus such that each G_{i+1} is derived from a node $A \in T_i$ (where T_i is the coinductive tree of G_i) and a clause C using a non-empty substitution θ_{i+1} . In this case, $\langle A, C, \theta_{i+1} \rangle$ is called a resolvent.

Coinductive derivations resemble *tree rewriting*. They produce the “lazy” corecursive effect: derivations are given by potentially infinite number of steps, where each individual step is executed in finite time.

Example 2.3 Figure 2 shows three possible coinductive derivations for **BitStream** and **BitList**. Note that two derivations for **BitList** terminate (with \square closing all branches). Note also, that this time, due to the and-or parallel nature of coinductive trees, changing the order of atoms or clauses in the program **BitList** will not change the result.

For terminating coinductive derivations, we require at least one or-subtree of the coinductive tree to be closed (with \square leaves). We also say in such cases that the coinductive tree contains a *success subtree*. The last coinductive trees in the second and third derivation of Figure 2 are themselves success subtrees.

Due to its and-or parallel properties (Komendantskaya et al. 2014c), CoALP is more robust than eager sequential SLD-resolution when it comes to reflecting program’s operational meaning; and mere change in the clause order would not place a terminating recursive function into a non-terminating class, cf. Figure 1. Yet more importantly, this new coinductive derivation procedure allows us to characterise productive and non-productive programs with better precision. In Introduction, we have seen that according to eager interpreter of CoLP, both programs **BadStream** and **TakeFirstN** are non-terminating; despite of one being productive, and another – non-productive. Next example shows that under lazy execution, productive programs with irregular pattern of resolvents can be handled more naturally.

Example 2.4 Figure 3 shows the first steps in the derivation for the program **TakeFirstN** and the goal $\text{taken}(s(s(0)), X)$. Unlike CoLP, CoALP is able to compute the second element of the stream in finite time.

There will be classes of non-terminating and non-productive programs for which coinductive trees grow infinite, and lazy derivations fail being “lazy”. The program **BadStream** is one such example. We will consider this issue in the next section.

3 Guarding (Co)recursion

The previous section introduced coinductive trees, which allowed us to distinguish terminating and productive programs like **BitStream**, **BitList**, **TakeFirstN** from non-productive programs like **BadStream**, by simply observing that coinductive trees remain finite for the former, while

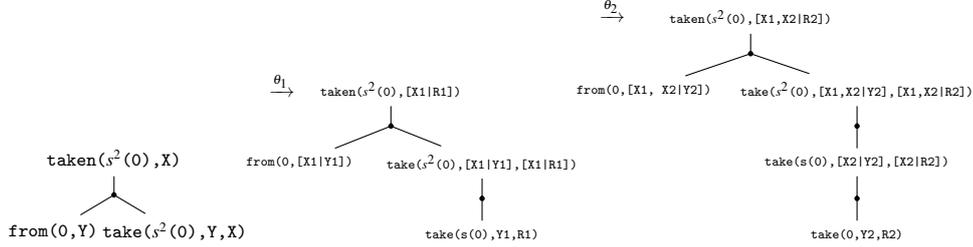


Fig. 3. First steps of the derivation for the goal $\text{taken}(s^2(0), X)$ – $s^2(0)$ denotes $s(s(0))$ – and the program *TakeFirstN*, with $\theta_1 = Y/[X1|Y1], X/[X1|R1]$ and $\theta_2 = Y1/[X2|Y2], R1/[X2|R2]$. As *take* is an inductive predicate, and *from* is coinductive; resolvents for *take* nodes are given priority.

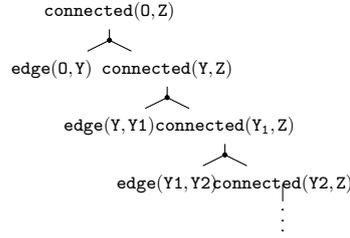


Fig. 4. The infinite coinductive tree for the program *GC* from Example 3.1, for the database $\text{edge}(0, 1) \leftarrow$.

growing infinite for the latter. It was especially significant that this new approach was, unlike Prolog, robust to permutations of clauses and atoms, and, unlike CoLP, was working with productive irregular streams. Curiously, the following logic program fails to produce finite coinductive trees:

Example 3.1 (GC) Let *GC* (for graph connectivity) denote the logic program

1. $\text{connected}(X, X) \leftarrow$
2. $\text{connected}(X, Y) \leftarrow \text{edge}(X, Z), \text{connected}(Z, Y)$

It would be used with database of graph edges, like $\text{edge}(0, 1) \leftarrow$.

The program gives rise to infinite coinductive trees, see Figure 4. It would terminate in LP, but, similarly to our discussion of **BitList**, would lose the termination property if the order of clauses (1) and (2) changes, or if the order of the atoms in clause (2) changes.

The reason behind infinity of coinductive trees for the above program is the absence of function symbols – “constructors” in the clause heads. The lazy nature of coinductive trees was in part due to the term-matching used to compute them. Term-matching loses its restrictive power in the absence of constructors. A very similar procedure of guarding recursion by constructors of types is used in e.g. Coq or Agda. This observation would suggest an easy way to fix the **GC** example, by introducing reducing dummy-constructors:

Example 3.2 (Guarded GC)

1. $\text{connected}(X, \text{cons}(Y, Z)) \leftarrow \text{edge}(X, Y), \text{connected}(Y, Z)$
2. $\text{connected}(X, \text{nil}) \leftarrow$

Considerations of this kind led us to believe that our lazy (co)recursive approach opens the way for a compile-time termination and productivity checks akin to respective checks in Coq or Agda (Coquand 1994; Abel et al. 2013). The programmer would be warned of non-terminating cases and asked to find a guarded reformulation for his functions. In Coq and Agda, different checks are imposed on recursive functions (“structural recursion” condition) and corecursive functions (“guardedness” checks). In logic programming terms, where types or predicate annotations are unavailable, we can formulate a uniform productivity property for recursive and corecursive programs, as follows:

Definition 3.1 *Let P be a logic program, P is productive if for any goal G , the coinductive tree for $P \cup \{G\}$ has a finite size.*

The above is a semantic property; syntactically, we need to introduce guardedness checks to ensure productivity. The intuitive idea is to ensure that every coinductive program behaves like `BitStream`: `BitStream` is guarded by the coinductive function symbol (or “guard”) `scons` (denoted by `[. | .]`); and hence all coinductive trees for it are finite, see Figure 2. On the contrary, **Comember** lacks a guarding constructor.

Example 3.3 (Comember) *The predicate `comember` is true if and only if the element X occurs an infinite number of times in the stream S .*

$$\begin{aligned} 1. \text{drop}(H, [H|T], T) &\leftarrow \\ 2. \text{drop}(H, [H1|T], T1) &\leftarrow \text{drop}(H, T, T1) \\ 3. \text{comember}(X, S) &\leftarrow \text{drop}(X, S, S1), \text{comember}(X, S1) \end{aligned}$$

Comember is un-productive for e.g. the coinductive tree arising from the query `comember(X, S)` contains a chain of alternating \bullet 's and atoms `comember(X, S1)`, `comember(X, S2)`, etcetera, yielding an infinite coinductive tree.

We will give a high-level formulation of guardedness checks here, for more technical discussion, see (Komendantskaya et al. 2014a).

Guardedness Check 1 (GC1): If the same predicate Q occurs in the head and in the body of a clause, then there must exist a function symbol f occurring among the arguments of Q ; such that the number of its occurrences is reduced from head to body.

Example 3.4 (Guarded Comember) *We propose the following guarded definition of `comember`; thereby simplifying it and reducing an extra argument to `drop`.*

$$\begin{aligned} 1. \text{drop}(H, [H|T]) &\leftarrow \\ 2. \text{drop}(X, [H|T]) &\leftarrow \text{drop}(X, T) \\ 3. \text{gcomember}(X, [H|T]) &\leftarrow \text{drop}(X, [H|T]), \text{gcomember}(X, T) \end{aligned}$$

In CoALP, the goal `gcomember(0, nats)` will lazily search for 0 in an infinite stream of natural numbers, but it never falls into un-productive coinductive trees, as CoLP would do.

GC1 would be sufficient for some programs, like `BitStream`, where there is only one (co)inductive clause; but not in the general case. LP in general is not compositional, that is, composing two programs may yield a program that has semantic properties not present in the initial programs.

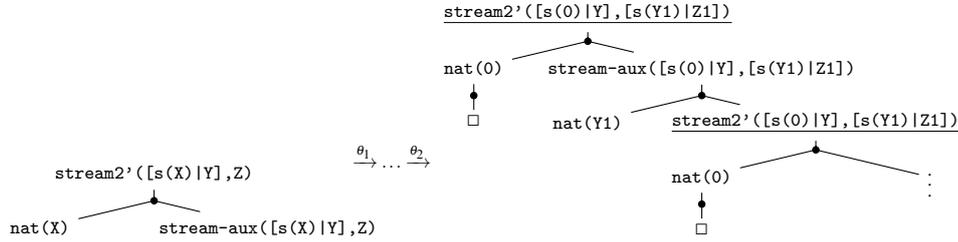


Fig. 5. Coinductive derivation of $\text{stream2}'([s(X)|Y], Z)$ and the program from Example 3.5 producing an infinite coinductive tree, with $\theta_1 = X/0$ and $\theta_2 = Z/[s(Y1)|Z1]$. The figure also represents one GC-derivation generated during GC3. GC3 detects the un-guarded loop; see the underlined atoms.

Same rule applies in CoALP: if both P_1 and P_2 are productive programs, their composition is not guaranteed to be a productive program; the next check is imposed to cover the compositional cases.

Guardedness Check 2 (GC2): For every clause head A , construct a coinductive tree with the root A . If there are atoms $Q(\bar{t})$ and $Q(\bar{t}')$ in the coinductive tree such that $Q(\bar{t}')$ is a child of $Q(\bar{t})$, apply GC1 to the clause $Q(\bar{t}) \leftarrow Q(\bar{t}')$.

GC1–GC2 handle some programs well, but they are still insufficient in the general case. The following program passes the checks GC1–GC2, but is not productive in the sense of Definition 3.1, see Figure 5.

Example 3.5 (Un-productive Program that passes GC1–GC2)

$$\begin{aligned} 1. & \text{stream2}'([s(X)|Y], Z) \leftarrow \text{nat}(X), \text{stream-aux}([s(X)|Y], Z) \\ 2. & \text{stream-aux}(X, [s(Y)|Z]) \leftarrow \text{nat}(Y), \text{stream2}'(X, [s(Y)|Z]) \end{aligned}$$

Guardedness Check 3 (GC3): For every clause head A , start a coinductive derivation with the goal A imposing GC2 condition to every coinductive tree in the derivation, and imposing the following termination conditions:

1. Terminate coinductive derivation if GC2 fails for at least one tree.
2. Terminate coinductive derivation if all branches are either closed with \square or contain guarded loops only.

Note that the checks GC1–GC3 we have introduced here are a pre-processing (compile-time) mechanism of CoALP. Once the program passed the guardedness checks, it does coinductive derivations lazily; and does not require any loop-detection procedures at run-time. If a program fails GC1–GC3, the programmer will be asked to re-formulate the definitions as we have seen in Examples 3.2 and 3.4. The first implementation of guardedness checks is available at (Komendantskaya et al. 2014b).

We finish this section with Table 1 comparing how SWI-Prolog, CoLP and CoALP handle various recursive and corecursive programs. For CoALP, we also benchmark guardedness checks.

For coinductive programs, CoLP can only handle coinductive programs that contain a regular pattern and fails otherwise (cf. Table 1); on the contrary, CoALP, in its lazy style, works for any program. This is illustrated, for instance, with the programs TakeFirstN and TakeRepeat. Table 1 shows that CoALP is slower than the CoLP interpreter and SWI-Prolog – note that SWI-Prolog is a fully-tuned mature programming language and the CoLP interpreter runs on top of SWI-Prolog, as opposed to our implementation of CoALP in Go from scratch.

		CoALP	CoLP	SWI-Prolog
TakeFirstN†	Yes	GC time: 0.0002s runtime: lazy execution	No	No
Takerepeat†	Yes	GC time: 0.0009s runtime: lazy execution	Yes (0.0001s)	No
Comember†		Not guarded	Yes [?] (0.0001s)	No
GComember†	Yes	GC time: 0.0011s runtime: lazy execution	Yes [?] (0.0001s)	No
SumFirstn†	Yes	GC time: 0.0013s runtime: lazy execution	No	No
FibStream†	Yes	GC time: 0.0006s runtime: lazy execution	No	No
Infinite Automata†	Yes	GC time: 0.0011s runtime: lazy execution	Yes (0.0001s)	No
Knights	Yes	GC time: 0.225s runtime: 3.002s	Yes (1.13s)	Yes (0.012s)
Finite Automata	Yes	GC time: 0.0011s runtime: 0.0023s	Yes (0.04s)	Yes (0.0005s)
Ackermann	Yes	GC time: 0.001s runtime: 13.23s	Yes (7.692s)	Yes (3.192s)

Table 1. Execution of different programs in CoALP, CoLP and SWI-Prolog. Examples marked with † involve both inductive and coinductive predicates. In the table, “No” means that the system runs forever without returning an answer, and “Yes[?]” indicates that the program succeeds if it contains a regular pattern and fails otherwise.

4 Work-in-Progress: Soundness for Corecursion

There are two main directions for CoALP’s development, both related to soundness:

(I) We are in the process of establishing soundness of **GC1-GC3** that is, the property that, *if a program P is guarded by **GC1-GC3**, then it is productive in CoALP.*

Proving this property in the general case is a challenge; and involves pattern analysis on re-solvents and also a proof of termination of **GC1-GC3**. Example 3.5 and Figure 5 give a flavor of the complicated cases the guardedness checks need to cover. Note that **GC1-GC3** provide the guarding property only in the CoALP setting, and the same idea of guarding (co)recursion by constructors would fail for standard LP or CoLP, as many examples of this paper show.

(II) Soundness of coinductive derivations needs to be established. This challenge is best illustrated by the following example.

Example 4.1 (Soundness for Comember) *To check the validity of a query in Comember (Example 3.3) for an arbitrary stream, one needs to satisfy two conditions: 1) finding an element to drop in a finite time, 2) finding guarantees that this finite computation will be repeated an infinite number of times for the given stream. CoLP would handle such a case for all streams that consist of a regular finite repeating pattern and will not be able to handle cases when the input stream is not regular. CoLP would fail to derive true or falsity of e.g. the query `comember(0, nats)`, where `nats` is the stream of natural numbers, as CoLP falls into an infinite non-terminating computation and fails to produce any response to the query. CoALP in its current implementation will handle any case of corecursion, including `comember(0, nats)`, but in its lazy, and therefore partial, style.*

Similarly, TakeFirstN falls into an infinite loop with CoLP, but unfolds lazily with CoALP, see Figure 3. Laziness on its own, however, does not guarantee soundness.

For inductive programs and recursive functions, CoALP yields the same theorems of soundness and completeness as classical LP (Lloyd 1987); cf. (Komendantskaya et al. 2014a). The

only adaptation to the already described coinductive derivation procedure is the requirement that the derivation terminates and gives an answer whenever a *success subtree* is found. Thus, generalisation of standard soundness and completeness for induction in CoALP is not very surprising.

Soundness of CoALP for coinductive programs is conceptually more interesting: it has to include a number of guarantees that need to be checked at compile-time and run-time, that is:

1. Identification of the guarding pattern coming from sound guardedness checks;
2. Guarantee that the guarding pattern will be produced in a finite number of derivation steps;
3. Guarantee that the guarding pattern will be re-produced an infinite number of times.

Item 3. in particular may allow for a few different solutions. In its basic form, it can be a repeated regular pattern, as it is done in CoLP. In a more sophisticated form, it can cover irregular patterns, as long as more involved guarantees of infinite execution are provided, cf. Example 1.4 and Figure 3.

To conclude, we have described a new method to analyse termination and productivity of logic programs by means of lazy guarded coreursion in CoALP, as outlined in Figure 1. We advocated a new style of programming in LP, where the programmer is in charge of providing termination or productivity measures for (co)recursive programs at compile-time, as it is done in some other declarative languages with recursion and coreursion. Finally, we outlined the main directions towards establishing soundness results for CoALP outputs.

References

- ABEL, A. ET AL. 2013. Copatterns: programming infinite structures by observations. In *POPL'13*. ACM SIGPLAN Notices, vol. 48. 27–38.
- BERTOT, Y. AND KOMENDANTSKAYA, E. 2008. Inductive and coinductive components of corecursive functions in Coq. *ENTSC 203*, 5, 25–47.
- COQUAND, T. 1994. Infinite objects in type theory. In *TYPES'93*. LNCS, vol. 806. 62–78.
- DE SCHREYE, D. AND DECORTE, S. 1994. Termination of logic programs: the never-ending story. *J. of Logic Programming 19–20, Supplement 1*, 199–260. Special Issue: Ten Years of Logic Programming.
- GUPTA, G. ET AL. 2007. Coinductive logic programming and its applications. In *ICLP'07*. LNCS, vol. 4670. 27–44. Interpreter Available at <http://www.utdallas.edu/~gupta/meta.html>.
- GUPTA, G. AND COSTA, V. 1994. Optimal implementation of and-or parallel prolog. In *PARLE'92*. 71–92.
- KOMENDANTSKAYA, E. ET AL. 2014a. Coalgebraic logic programming: from semantics to implementation. *J. Logic and Computation*.
- KOMENDANTSKAYA, E. ET AL. 2014b. CoALP webpage: software and supporting documentation. <http://staff.computing.dundee.ac.uk/katya/CoALP/>.
- KOMENDANTSKAYA, E. ET AL. 2014c. Exploiting parallelism in coalgebraic logic programming. *ENTCS 303*, 121–148.
- KOMENDANTSKAYA, E. AND POWER, J. 2011. Coalgebraic derivations in logic programming. In *CSL'11. LIPICs*. Schloss Dagstuhl, 352–366.
- LLOYD, J. 1987. *Foundations of Logic Programming*, 2nd ed. Springer-Verlag.
- SIMON, L. ET AL. 2007. Co-logic programming: Extending logic programming with coinduction. In *ICALP'07*. LNCS, vol. 4596. 472–483.