



HAL
open science

Logic Programming Environments: Dynamic Program Analysis and Debugging

Mireille Ducassé, Jacques Noyé

► **To cite this version:**

Mireille Ducassé, Jacques Noyé. Logic Programming Environments: Dynamic Program Analysis and Debugging. [Research Report] RR-2618, INRIA. 1995. inria-00074067

HAL Id: inria-00074067

<https://inria.hal.science/inria-00074067>

Submitted on 24 May 2006

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

INSTITUT NATIONAL DE RECHERCHE EN INFORMATIQUE ET EN AUTOMATIQUE

***Logic Programming Environments:
Dynamic Program Analysis and Debugging***

Mireille Ducassé and Jacques Noyé

N° 2618

July 1995

PROGRAMME 2



*R*apport
de recherche



Logic Programming Environments: Dynamic Program Analysis and Debugging

Mireille Ducassé * and Jacques Noyé **

Programme 2 — Calcul symbolique, programmation et génie logiciel
Projet LANDE

Rapport de recherche n° 2618 — July 1995 — 43 pages

Abstract: Programming environments are essential for the acceptance of programming languages. This survey emphasizes that program analysis, both static and dynamic, is the central issue of programming environments. Because their clean semantics makes powerful analysis possible, logic programming languages have an indisputable asset in the long term.

This survey is focused on logic program analysis and debugging. The large number of references provided show that the field, though maybe scattered, is active. A unifying framework is given which separates environment tools into extraction, analysis, and visualization. It facilitates the analysis of existing tools and should give some guide lines to develop new ones.

Achievements in logic programming are listed; some techniques developed for other languages are pointed out, and some trends for further research are drawn. Among the main achievements are algorithmic debugging, tracing for sequential Prolog, and abstract interpretation. The main missing techniques are slicing, test case generation, and program mutation. The perspectives we see are integration, evaluation and, above all, automated static and dynamic analysis.

Key-words: Programming environments, logic programming, program analysis, debugging, survey article.

(Résumé : *tsvp*)

Paru dans *Journal of Logic Programming*, vol. 19/20, May/July 1994, pp 351-384. Anniversary issue: Ten years of Logic Programming.

*IRISA/INSA : 20 av des Buttes de Coësmes, 35043 Rennes ; email : Mireille.Ducasse@irisa.fr

**IRISA/INRIA : Campus Universitaire de Beaulieu, 35042 Rennes ; email : Jacques.Noye@irisa.fr

Environnements de Programmation Logique : Analyse Dynamique de Programme et Débogage

Résumé : Les environnements de programmation sont essentiels pour l'acceptation des langages de programmation. Cet article de synthèse souligne que l'analyse de programme, tant statique que dynamique, est le point central des environnements de programmation. Grâce à leur sémantique "propre" qui permet des analyses puissantes, les langages de programmation logique ont un avantage indiscutable à long terme.

Cette synthèse est centrée sur l'analyse et le débogage de programmes logiques. Le grand nombre de références bibliographiques atteste que le champ, bien que dispersé, est actif. Nous présentons un cadre unificateur qui sépare les outils d'environnement en trois phases : extraction, analyse et visualisation. Ce cadre facilite l'analyse des outils existants et donne des repères pour en développer de nouveaux.

Les accomplissements de la programmation logique sont listés ; quelques techniques développées pour d'autres langages sont mentionnées ; enfin, quelques directions de recherche sont proposées. Parmi les accomplissements principaux de la programmation logique on trouve : "algorithmic debugging", les traces de programmes Prolog séquentiel et l'interprétation abstraite. Les techniques manquant principalement sont : "slicing", la génération de données de test et "program mutation". Les perspectives que nous voyons sont : l'intégration, l'évaluation et, surtout, l'analyse automatisée de programme, tant statique que dynamique.

Mots-clé : Environnements de programmation, programmation logique, analyse de programme, article de synthèse.

1 Introduction

Programming environments are essential for the acceptance of programming languages. Programmers, if not provided with the support of proper development and maintenance tools, are more likely to waste time and produce low quality software. They are therefore reluctant to use a language without appropriate programming environments, however powerful the programming language. Logic programming languages are no exception.

Logic programming is continually used as a basis for new languages; people need support to get acquainted with them. Furthermore, and somehow paradoxically, it is the very power of the declarativeness of logic programming which demands powerful support. The discrepancies between the declarative and operational semantics of logic programming languages can be a source of confusion for programmers; tools are required to bridge the gap.

Although logic programming requires programming environments like any other programming language, it offers a better basis than many of them. Clean semantics pave the way for powerful program analysis, which, as discussed in the following, is the core of programming environment tools. Even if not always developed with programming environment considerations in mind, logic program analysis is a flourishing activity. There is undoubtedly a breakthrough potential for programming environments in general.

Research on logic programming environments started relatively early [104] and shows substantial achievements. Original tools are now being adapted to other languages, for example Algorithmic Program Debugging [168] is being adapted to imperative languages [100] (see section 6.1). On the other hand, not all the techniques existing for other languages have been adapted to logic programming. The most striking example is a dependency analysis technique used for debugging, called “slicing” [192] (see section 6.3).

It has been emphasized that, in general, more research on programming environments is needed [148]. This is especially true in logic programming, where research on programming environments has been marginal in comparison to the needs. We believe that programming environments should be one of the top priorities for the logic programming community in the next ten years.

This survey gives a unifying framework which should help to develop further research on the topic. It is subjective and certainly incomplete, but should provide a useful starting point for people interested in logic programming environments. It gives an overview of techniques, tools, and systems which help programmers develop and maintain logic programs. It goes into some details for very basic techniques. Wherever it seems necessary, it mentions techniques developed for other languages, which could be useful for logic programming.

The article is organized as follows. We first set the scene. In particular, we isolate the three aspects we believe are the core of programming environments: *program development*, *program analysis*, and *program debugging*. We then concentrate our discussions on program analysis and program debugging. Program analysis is especially emphasized as it is the basis of all the tools. Presentation and visualization of information is also addressed. Lastly, we discuss two general issues: integration mechanisms and evaluation criteria.

2 Setting up the scene

2.1 A user viewpoint of programming environments

The notion of “programming environment” can be understood in many different ways. A common view considers all aspects related to programming, such as language features, compiler, system support, tools, etc, as part of the programming environment. This view is, for example, illustrated by the wide range of topics addressed in a recent workshop about programming environments [102].

Although this view has some foundations, it is too broad to efficiently gather research. We would like to cover a more focused area. In particular, we want to distinguish research on programming environments from research on languages.

Research on programming environments can contribute significantly to the design of new languages, mainly by pointing out the weaknesses of existing programming languages. A good example is contextual logic programming, which is the main outcome of the Esprit project ALPES on logic programming environments [3]. Contextual programming neatly integrates modules inside Prolog [134]. The types, modules, and meta-programming facilities provided by the declarative language Gödel also owes a lot to programming environment concerns [90]. Another example is λ Prolog [139], which, among other features, provides “for free” the programming cliché mechanisms described in [7].

There are, however, problems which are not due to weaknesses of programming languages but to intrinsic weaknesses of the human beings who specify, develop, and maintain programs. We consider that excising language weaknesses is part of language research and that research on environments should concentrate on *tools to accommodate user weaknesses*. Thus, logic programming environments would still be needed even if the “perfect” logic programming language could be defined.

In particular, *development, analysis, and debugging* will always have to be supported. Firstly, some support is needed to help users develop their programs because this is a creative process. There is very little chance that we can ever fully model it. Secondly, automated analysis is needed because, however abstract the programming paradigm, the size of the programs will always end up too large for the human mind. Programmers will always have difficulties in understanding large programs [36]. Lastly, debugging and testing support is needed because, however careful programmers are, they can always make mistakes, either at the specification or programming level.

The classification into development, analysis, and debugging is mainly a guide line to present the survey. The boundaries between the three classes are not clear cut; for example, development and debugging are intertwined, and debugging requires program analysis. In the following we will concentrate on program analysis and debugging.

2.2 User-oriented program analysis

Program analysis is often performed in the context of compilation, aiming at optimization. When performed in the context of programming environments, the main difference is that

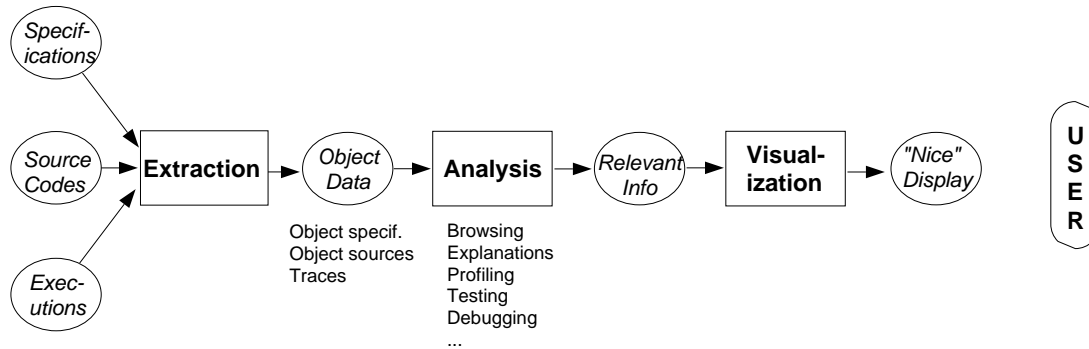


Figure 1: A model of user-oriented program analysis

the derived properties are ultimately analysed by human beings. This induces strong requirements about flexibility, response time, and ergonomy. It therefore influences the technical design of the related tools.

User-oriented program analysis is mainly aimed at helping programmers understand programs and their behaviours. They are usually referred to as debugging tools. They, of course, do help programmers to debug because one of the main problems while debugging is understanding programs. However, on the one hand, debugging is more than just understanding and, on the other hand, understanding can be an aim per se. For example, a mode analysis can help with the understanding of the dataflow of a program; a visualization tool can help with the understanding of the operational semantics of a language; when a new feature has to be added in an existing program the programmer has to understand this program, at least partially, to avoid introducing inconsistencies.

Figure 1 gives a model of user-oriented program analysis. This model is used as a guideline in the following sections, which describe each aspect in some depth. Let us just give a brief overview here. Program analysis uses three sources of data: program source codes, program executions, and program specifications. These data should be available to the analysis component as *object* data. This includes the transformation of specifications and source codes into an appropriate abstract syntax as well as the extraction of executions traces. The analysis component can perform browsing, explanation, profiling, testing, debugging, etc. It produces information, relevant to the user at a certain moment, which is displayed in an ergonomical way by a visualization component. This material is eventually examined by the user.

The model emphasizes two important aspects. Firstly, program analysis is not restricted to static analysis, i.e. analysis of source code only; section 3 stresses that specifications and executions are also interesting sources of data. Secondly, the central component of a programming environment is the *analysis component*.

In our view, the most important objective of a programming environment tool is to *automate as much as possible of the related development and maintenance task* in order to reduce the burden of the programmers. These tasks are usually too difficult to be fully automated. The visualization module is therefore needed to report to the user, but it is not the main part of the environment. We, as users, prefer to examine three lines of *relevant* information than a screen full of *nice*, graphically displayed, irrelevant information. Hence the analysis component should be the main focus of the research on programming environments.

Very often, existing tools are regarded too low-level or even useless because their analysis component is, actually, too weak. For example, tracers provide a limited number of filtering functionalities which are a degenerate and “hacked” form of analysis. Visualization tools usually provide more sophisticated analysis mechanisms to enable several levels of abstraction, but these mechanisms are still hidden, ad hoc and intertwined with the graphical mechanisms. With a proper analysis component the generality of these tools could be greatly improved.

The model is simple, if not simplistic. In particular, users also input requests to the tools. It should be stressed, however, that existing systems are only a degenerate form of this model. No current system satisfactorily covers extraction, analysis, and visualization. We hope that in the near future we will see systems covering and integrating all aspects.

Note that program analysis as considered here is essentially a meta-programming activity; it handles program as data, as well as data as programs. This suggests considering the analysis tools (if not the whole environment) as a meta logic program and the target logic language as the implementation language. Some important benefits can be expected from this approach. In particular, the development of new tools is facilitated by the self-applicability of the environment and the possibility to reuse software from the object level. This also provides a demanding testbed for the meta-programming facilities of logic programming languages.

3 Program analysis data

As already mentioned, program analysis can use three sources of data: *program source codes*, *program executions*, and *program specifications*. Program source code is a natural candidate; there are indeed many projects on static analysis. In this section we want to clarify the role of program executions and program specifications.

3.1 Program executions

Program analysis seems often restricted to static analysis, i.e. analysis of program source code, only. Although static analysis has advantages, it is not always the most appropriate type of analysis.

As shown by the popularity of prototyping, it seems, for example, that people often have a better understanding of what a program should *do* than of what it should *be*. As

understanding is a central issue in user-oriented program analysis, executions which show what programs actually do are an essential source of data.

As a matter of fact, there exist a number of dynamic analysis tools, for example, debuggers, profilers, monitors, or explanation generators. These tools are often ad hoc, with limited scope. They are, however, driven by user demand, and their usefulness is not to be questioned.

An execution is an instantiation of a program with specific input; only some of the potential paths of a program are covered. While this reduces the generality of the analysis, it also reduces its complexity. In some cases, such as bug diagnosis, there is no need to investigate all potential execution paths. Only the paths actually used to produce the error symptoms are relevant. Hence, analysing executions is sometimes more relevant and more efficient.

This is not to say that dynamic analysis is better than static analysis. Sometimes one is more appropriate, sometimes it is the other one. Our experience shows that, most of the time, a combination of both is actually the most appropriate.

3.2 Specifications

The third source of data for program analysis is program specifications, where a program specification is, basically, a high-level description of the *intention* of the programmer.

A big asset of logic programming is that it provides a smooth transition between specifications and programs. In many cases, specifications can be made formal and executable [82], i.e. the distinction between specifications and programs is mainly a question of efficiency and level of abstraction. As a matter of fact, Prolog is sometimes considered as a specification language (see for example [13, 129, 93]). Even when, for the sake of expressiveness, specifications cannot be reduced to logic programs, a uniform use of logic provides a path between specifications, (pure) logic programs, and (efficient) Prolog programs. This has been applied with success to program synthesis (see the survey on program synthesis, by Deville and Lau, in this issue [49]).

The link between specifications and programs is currently tightened through a growing understanding of the importance of types [153], which are essential to capture the notion of *meaningless* or *erroneous* data as well as to reveal meaningful approximations of the relations to be computed. This has led to the introduction of typed versions of Prolog, e.g. Typed Prolog [110], as well as languages which further depart from Prolog, e.g. Gödel [90] or λ Prolog [139]. The type systems offered by these languages may not, however, be expressive enough to capture the intention of the programmer. Naish suggests, for instance, to include distinguished predicates (called *type predicates*) at the specification level. These predicates are not part of the program level but can be used to make more precise the intended semantics of the program in different phases of its development [142].

Introducing specifications replaces a big step (programmer's intention to program) by two smaller steps (intention to specification and specification to program), hopefully easier to manage. Capturing the programmer's intention in complete and sound specifications remains a difficult task. In the best case, methodologies are used to enforce rigorous development

based on specifications; in the worst case, whether specifications are produced or not depends entirely on the good will of programmers.

As a result, program analysis cannot rely on the availability of *complete* specifications. Nevertheless, there are always some aspects of a system which can be easily specified. Thus, expecting *incomplete* specifications is reasonable and may be beneficial. For example, assertions partially specifying the behavior of a program help reduce user's mediation during algorithmic debugging [56].

4 Extraction

Program specifications, program source codes, and program executions cannot be straightforwardly used as analysis data. The *extraction component* transforms them into object data. The role of the object syntax is both to make sure that there will be no confusion between the object program (the program being analyzed) and the meta program (the analyzer), and to retain information which may be pertinent to the user (e.g. variable print names).

The transformation of program sources and specifications into object data are simple syntactic transformations. Information about program executions is not as straightforwardly available. Program executions are modelled by sets of events, usually called *traces*, or *event histories*. Instrumentation depending on the chosen execution model is required to extract these traces.

In the remainder of the section we briefly describe each of these mechanisms. They are basic mechanisms for general static and dynamic analysis, essential to programming environments.

4.1 Object data

The basic role of the transformation into object data is to avoid any confusion between the object-level and the meta-level. It is essentially a naming problem, well known in static analysis and, more generally, in meta-programming [124, 8, 121, 196].

4.1.1 Non-ground and ground representations

The major issue is the representation of object-level variables. The general alternative is between a non-ground and a ground representation. Object-level variables are represented by meta-level variables in non-ground representations and by ground terms in ground representations.

Ideally, a non-ground representation should be typed as advocated by Hill and Lloyd [91] in order to prevent unifying real meta-level variables with meta-level variables representing object-level variables.

Using typeless Prolog, a straightforward choice is to use a ground representation, which has the advantage that the meta-level version of the meta-logical predicates can be given a declarative semantics [91].

A basic ground representation of terms is described by Barklund in [8] as follows:

- The name of a constant \mathbf{X} is the compound term `constant(\mathbf{X})`.
- The name of a compound term whose functor is \mathbf{F}/\mathbf{m} and whose arguments are $\mathbf{A1}$ to \mathbf{Am} is `compound(\mathbf{F} , [$\mathbf{N1}$, ..., \mathbf{Nm}])`, where $\mathbf{N1}$ to \mathbf{Nm} are names of the arguments $\mathbf{A1}$ to \mathbf{Am} , respectively.
- A name of a variable \mathbf{X} is a term `var(\mathbf{I})` where \mathbf{I} is any positive integer, such that the same integer \mathbf{I} is used for every occurrence of \mathbf{X} in all terms under consideration, and that all distinct variables in these terms are assigned different names.

Many variants are possible. Actually, van Harmelen advocates user definable representations [187]. For our purpose, the name of constants can be the constants themselves. It is also possible to represent compound terms by compound terms with the same principal functor. It is then necessary to be careful not to confuse the representation of a compound term `var(\mathbf{Term})` and the representation of a variable. A solution consists of applying a mechanism reminiscent of the double quote mechanism used to insert quotes in strings and represent the compound term as `var(var(\mathbf{Term}))`. Finally, the representation can also be supported at the implementation level, e.g. by replacing wrappers such as `var/1` by special tags [8].

The use of such a representation is classical in static analysis and meta-programming. However, to our knowledge only interpreter-based tracers, such as Maeda's [127], have extended its use to the representation of events.

4.1.2 Unique (global) identifiers

Traced events including variables should not lose variable sharing information. If two events involve the same variable, it is necessary for dataflow analysis that the analysis component can recognize it. For instance, in Figure 2, one expects to see the same name for the first argument of `p/1`, line 1, and the argument of `q/1`, line 4. Basically, each new variable created during an execution should be associated with an identifier unique within the scope of the execution subject to analysis. In the terms of the ground representation defined above, all terms under consideration for the definition of \mathbf{I} are all the terms involved in the events susceptible to analysis. However, since dataflow analysis only makes sense along a given derivation, distinct variables which are not shared between different derivation paths may be assigned the same name. Let us look again at Figure 2. It does not make sense to look at lines 2 and 9, where `_2 = a` and `_2 = b`, respectively, without understanding that they belong to two different derivations (one from line 2 to 7, and one from line 9 to 12). Therefore, the two different instances of \mathbf{X} , lines 5 and 12, can share the same identifier without creating any ambiguity.

The standard practice is to use the physical address of a variable as its identifier. Unfortunately, the uniqueness of such identifiers cannot be guaranteed; addresses may change during the execution (e.g. due to a globalisation in a WAM-based implementation, or a memory copying or compaction due to garbage collection).

```

p(X, Z) :- Z = a, q(X), fail.          q(f(X)).
p(X, Z) :- Z = b, q(X).

?- p(X, Y).

1      Call   p(_1, _2)
2      Call   _2 = a
3      Exit   a = a
4      Call   q(_1)
5      Exit   q(f(_3))
6      Call   fail
7      Fail   fail
8      Next   p(_1, _2)
9      Call   _2 = b
10     Exit   b = b
11     Call   q(_1)
12     Exit   q(f(_3))
13     Exit   p(f(_3), b)

```

Figure 2: A simple Prolog trace with variable identifiers

4.1.3 Variable print names

Usually, programmers carefully choose the *print* names of program variables (i.e. their external, source-level, representation), so that they carry some meaning about the application domain. As the result of an analysis is to be presented to users, it is essential that the variable print names are kept in the representation of object variables. The representation of a variable should therefore be composed of both its print name and its unique global identifier. Prolog/KR [144], KCM [146], and SEPIA [131] have implemented such a representation (modulo the problem mentioned above).

4.2 Basic execution models

Information about executions is not straightforwardly available. Basic models of executions have to be designed so that extraction mechanisms can extract information related to these models. It should be emphasized again that the basic models are not aimed at users, but at the automated analysis module.

An execution is modelled by a set of atomic actions, or *events*, organized as a *trace*, or *event history*. A sequential execution leads then to a single sequence of events, whereas a parallel execution leads to a set of (local) sequences of events connected by dependency links.

Events can potentially be very low level, e.g. the execution of a machine instruction, a store access, or the execution of a communication primitive. The issue is to design a model which is sufficiently detailed so that the analysis does not miss any information, sufficiently high-level so that the analysis does not spend too much time restructuring or recomputing information, and sufficiently easy to extract so that the resulting response times of the tools are reasonable.

For Prolog, a number of basic models exist. A popular one is the “box model”, designed by Byrd [26] and described formally by Tobermann and Beckstein [182]. This model shows a procedure execution as a black box (hence its name). Events are related to goals; there are four types of events (traditionally called “ports”). For a given goal g , the meaning of the events is as follows:

Call g : g is being invoked;

Exit g : g has just been proved;

Redo g : the execution backtracks to a *subgoal of g* ;

Fail g : g could not be proved.

Another model is closer to the operational semantics of Prolog. It shows backtracking in the order in which it actually appears in the execution. The only difference with the box model is the backtracking port **Next**, which has a slightly different semantics than **Redo**¹:

Next g : the execution backtracks to g .

The difference between the two models is subtle but significant. Let us illustrate it on an example. Figure 3 shows a program traced using the five ports previously defined. One can see that for the facts (e.g. **s**) the **Redo** and **Next** events are consecutive and seem redundant. But in large executions, goals with a non empty body are usually more numerous than facts. If such goals have choice points (e.g. **q**), the two backtracking events related to these goals occur at very different places. Here, re-entering the box of **q** is shown at line 7 (before the very first backtracking to a subgoal of **q**), while the actual backtracking to **q** is shown at line 11. The deeper the subtree to prove goals such as **q**, the larger the distance. Note that if the line 11 is missing from the trace (**Next q**), which is the case in the box model, the invocation of **t** on line 12 seems to arrive by magic.

The two ports are useful. The **Next** port is more faithful to the operational semantics of Prolog, it tells where backtracking actually occurs. The **Redo** port is very useful to trace “breadth-first”. In particular, if the details of the execution of a goal are hidden, the **Redo** port captures backtracking information which otherwise would have been lost [188]. For automated analysis both ports should be extracted and clearly separated as in the example.

Many refinements to the previous models can be made. For example, Eisenstadt [69] defines six types of failure: “subgoals of g have failed; system primitive has failed; subgoal

¹Note that the two ports are often given the same name in different models, which is a source of confusion.

```

q :- s.    q :- t.    s.

?- q, r.

1    Call q
2    Call s
3    Exit s
4    Exit q
5    Call r
6    Fail r
7    Redo q      % *** re-entering the box of q
8    Redo s      % re-entering the box of s
9    Next s      % backtracking to s
10   Fail s
11   Next q      % *** backtracking to q
12   Call t
13   Fail t
14   Fail q

```

Figure 3: An operational model with two types of backtracking events

has backtracked to cut, failing parent; no definition in database; definition exists but different arity; arity OK but no resolvents found; and clause would have been considered but does not unify”. Morishita and Numao [135], as well as Schleiermacher et al. [164], add information at the clause level, in particular about unification. The box model has been adapted to functional logic programming by Hanus and Josepfs [88].

For some analyses, e.g. algorithmic debugging, the control flow information is of little interest and a declarative model is more appropriate than an operational one. In such a case, the resulting proof tree can be sufficient. Sometimes, it is also important to have information about intermediate proof trees (i.e. before failures) or to see how the resulting proof tree has been constructed. In such cases one can use a degenerate version of the previous operational models and keep only the **Exit** ports.

This list is certainly not exhaustive and each Prolog tracer actually uses a different variant of the previous models with more or less information. Each model has its advantages, but none of them is absolutely better than all the others. The basic model to be used depends on many parameters, in particular on the type of analysis which is to be performed.

4.3 Tracing mechanisms for sequential executions

In the following, we list and briefly discuss some mechanisms to extract trace information from sequential executions. A more thorough discussion can be found in [67].

4.3.1 Manual program source instrumentation

A straightforward way to extract information about executions of a program is to instrument its source code. The most primitive way is to let users insert “write” statements into their programs. This can be very precise when users know exactly what information is needed and insert write statements at appropriate places, but this manual treatment can become very tedious. It usually requires a lot of trials and errors to find the relevant places.

A more sophisticated way consists of inserting (still by hand) “trace” statements instead of “write” statements. The trace statements can then be programmed to send information to the analysis component. Nevertheless, letting the user insert the trace statements is not appropriate for automated dynamic program analysis. The trace information has to be generated automatically so that the analysis component has some guarantees that the required information will be extracted.

4.3.2 Automatic program source instrumentation

The next step is to automatically instrument programs. A technique to trace according to the box model described in the previous section is to encapsulate examined predicates. For example, the definition of predicate `p/1` can be replaced by:

```
p(X) :-
    ( trace(call, p(X)) ; trace(fail, p(X)), fail ),
    p_do(X),
    ( trace(exit, p(X)) ; trace(redo, p(X)), fail ).

p_do(X) :-
    <p body>.
```

The `trace/2` predicate succeeds once. When `p/1` is invoked, `trace(call, p(X))` is invoked and a `Call` line is traced, then `p/1` is executed normally. If `p/1` succeeds, `trace(exit, p(X))` is invoked, otherwise `trace(fail, p(X))` is invoked and the failure is propagated. If `p/1` succeeds and the execution later backtracks, `trace(redo, p(X))` is first invoked, then `fail/0` forces a backtracking to `p/1` itself.

Note that it is the current substitution of the variables which is traced. This is in general what is needed; the `Redo` and `Fail` events, however, need some refinements. This mechanism is also more sophisticated for full Prolog ².

Note also that, as mentioned before, if one is only interested in the construction of the proof tree, tracing the `Exit` events is sufficient. The transformed `p/1` predicate is then:

```
p(X) :-
    p_do(X),
    trace(exit, p(X)).
```

²Interested readers can see the “advice” utility written by R. O’Keefe which is part of the DEC10 Prolog library, available by anonymous ftp from the AIAI of the University of Edinburgh (aiai.edinburgh.ac.uk).

The `trace/2` predicate is programmed to pass the information to the analysis component. The transformed code can be compiled, and therefore runs reasonably efficiently, although the method leaves choice points at each predicate even when it is determinate, impeding last call optimization. This solution is not too difficult to implement and is satisfactory for occasional breakpoints.

When only the resulting proof tree is of interest, another transformation consists of adding an extra argument to each predicate, which collects the proof tree. The predicates

```
p :- q, r.
q :- ( s1 ; s2 ).
```

are transformed into something like:

```
p(proof(p, (ProofQ, ProofR))) :- q(ProofQ), r(ProofR).
q(proof(q, ProofS)) :- ( s1(ProofS) ; s2(ProofS) ).
```

Note that this can only answer questions of the type “why X” for programs without negation. In order to answer questions of the type “why not X”, and to be able to trace programs with negation, failing search trees have to be explicitly stored [190, 172].

As discussed by Specht [172], the latter approach seems to be particularly adapted to tracing bottom up evaluations such as in deductive databases. Indeed, in bottom up evaluations one cannot rely on the execution order to trace information by side-effect. Furthermore, in deductive databases the program is typically heavily transformed before it is compiled. Such a program instrumentation allows the information to be traced in terms of the initial program and not in terms of the transformed one.

4.3.3 Instrumentation of meta-interpreters

A natural step is to use the same techniques as before but within a meta-interpreter. An introduction to meta-interpreters can be found, for example, in [174, Chapter 19]. Typically, a meta-interpreter contains some clauses to recursively process lists of goals, and a clause which reduces a single goal. This clause can be instrumented with the same tracing instructions as above.

```
solve(Goal) :-
  ( trace(call, Goal) ; trace(fail, Goal), fail ),
  reduce(Goal),
  ( trace(exit, Goal) ; trace(redo, Goal), fail ).
```

This can be generalized to multi-layered meta-interpreters [14, 175, 34].

In the same way, the construction of the proof tree in an argument, mentioned in the previous paragraph, can be done by a meta-interpreter [19, 167]. The `solve/2` predicate carries on a second argument which is constructed inside the `reduce/2` predicate as follows:

```
reduce(Goal, proof(Goal, ProofBody)) :-
  clause(Goal, Body),
  solve(Body, ProofBody).
```

Unfortunately, the space and time optimizations of the compiler are mostly lost. Thus, large programs that can run when compiled, may require too much memory when interpreted by a meta-interpreter. If the only tracing means is a meta-interpreter, programmers may be left with programs which they cannot trace. Furthermore, for user-oriented analysis, response times must be kept reasonable, hence the lack of efficiency is a major problem.

Partial evaluation [103, 84, 161, 173, 124], used on the instrumented meta-interpreter and the program to be traced, can produce a program equivalent to the instrumented program of the previous section. Partial evaluation plus instrumented meta-interpreters is more generic than a specialized program instrumentation. One does not need to write a specialized program transformation but can benefit from a general partial evaluator. Furthermore, if the simple instrumented meta-interpreter is efficient enough to trace the analysed program, then the partial evaluation step can be skipped.

Enhancing standard meta-interpreters to get high-level information is very easy. They are therefore useful to prototype particular analyses. For example, [174, chapter 19] presents a meta-interpreter which detects stack overflows, one which diagnoses false solutions bottom up, one which diagnoses false solutions top-down, and one which diagnoses missing solutions.

4.3.4 Compiled code instrumentation

The implementation of the tracer can be pushed further down. If an interpreter implemented in a low-level language is used, this interpreter can be modified to generate trace information. Alternatively, if a compiler is used, the compiled code can be modified for the same purpose. In the case of byte-code compilation, the emulator can also be modified.

Notification points are set into the compiled code when an interesting event is reached. Note that not every interesting event corresponds to a physical location in the compiled code. The most obvious example is the **Exit** port. All reasonable Prolog compilers implement last call optimization via a continuation mechanism; the last subgoal call of a clause does not return to its caller, but the execution directly proceeds with the continuation of the call. Hence there is no location in the code where the debugger could be notified when the last subgoal of a clause exits.

In the tracer of the ECLiPSe/Sepia Prolog system [131] the **Call** port corresponds directly to a call instruction, while only facts notify the debugger when they exit. The tracer uses its own stack of **Call** and **Exit** frames to reconstruct the other ports.

Compiled code instrumentation has given, so far, the best results, in terms of both efficiency and space consumption. It has also led to more precise tracing. For example, cuts can be traced showing precisely which choice points are actually removed. This solution, however, takes notably more time and effort to implement. The resulting tracer is also highly dependent on the implementation of the interpreter or compiler it is tightly connected with. It is therefore worth considering the other solutions further, all the more as program transformation and partial evaluation techniques have great potential.

4.4 Tracing mechanisms for parallel executions

Tracing parallel executions exhibits further difficulties due to the “probe effect”, nonrepeatability, and the lack of a synchronized global clock [130]. The first two problems can be alleviated by resorting to execution replay. The latter one requires a careful organization of the trace, capturing the partial ordering of the recorded events.

4.4.1 Execution replay

The probe effect and nonrepeatability are actually two aspects of the same problem: parallel executions inherently lead to *races*, i.e. competitions for the access of some shared data, for example, OR-parallel access to a choice point, or AND-parallel access to an unbound variable. The resolution of a race depends on the relative speeds of the competitors and is therefore sensitive to any factor modifying these relative speeds, e.g. the state of the memory hierarchy or the amount of network traffic.

This problem is especially acute with concurrent logic languages where indeterminism, through the committed choice mechanism, is a basic language feature. If several clauses match a goal, the chosen clause depends on the availability of the goal variable bindings and the relative speeds of the parallel clause matching activities. As a result, running a program twice can lead to different behaviours of the program. The problem is not absent from parallel Prolog systems either. In OR-parallel systems, such as Aurora [32], the relaxation of the order in which the solutions are produced, combined with the use of cuts or one-solution constructs, also leads to programs returning different results, here different answer substitutions, including failure, from one run to the other. The introduction of asynchronous built-in side effects (I/O, assert/retract) adds some more sources of indeterminism.

Repeatability remains an issue even when there is no indeterminism in the language (or in the program). In that case, the results are repeatable, but the paths followed to get to these results may still be different. Let us consider the AND-parallel execution of two goals for both of which there is no successful derivation. Failure may then be caught alternatively while solving one goal or the other. At an even lower level of observation, the paths may be the same, but the scheduling of work between the processing agents different. This cannot be ignored when tracing is done with performance debugging in mind or when the system itself is being debugged.

Finally, since tracing is both time and space consuming, switching tracing on may also influence race resolution and modify the behaviour of the traced program. This is the probe effect. In general, a trace does not reflect the execution without tracing and two executions of the same program do not result in the same trace.

Control driven execution replay [115, 118] is a generic method to deal with indeterminism in parallel programs. During an *initial execution*, each process involved in the computation records a minimal trace, called *history tape*, collecting the result of all the race resolutions, as well as program input. The race related events are very simple. In case of a shared memory implementation, each event corresponds to a shared object access. It records the object identifier and its version number. In case of a message passing implementation, each

event corresponds to a message reception. It records the sender identifier and the message number (local to the sender). This simplicity, together with the lack of global bottleneck, makes the process very lightweight, limiting the probe effect. The history tapes can then be used to guide new executions, called *replays*, which are equivalent to the initial execution. In particular, detailed tracing can be safely performed during replay.

Execution replay has been implemented in the parallel Prolog system PEPSys [9], based on Instant Replay [115], the seminal version of execution replay based on shared objects. Execution replay was mainly seen as a tool to help debug the system itself. In particular, it captured scheduling decisions. Interestingly, exploiting specific properties of the PEPSys computational model made it possible to implement a simplification of the initial instant replay scheme [44]. This resulted in a very low overhead for the initial execution, typically 2% to 3%, ensuring, in general, an identical behaviour of the system with and without recording on.

Execution replay has also been considered in the context of concurrent logic languages. In [167], Shapiro describes a simple Flat Concurrent Prolog meta-interpreter which, combined with a source transformation of a (non interactive) program, computes, as a tree data-structure, a trace recording the order of goal reductions. This makes it possible to reconstruct the computation, e.g. to apply algorithmic debugging techniques. In the same spirit, Gaifman et al. give a transformation which captures the indeterminism of committed choices [83]. The transformation is justified by a study of execution replay in an abstract setting, showing in particular that execution replay of concurrent logic/constraint programs is simplified by a monotonicity property of these programs. Preliminary experimental results on a stream merger running on the Concurrent Prolog system Logix [170] give some upper bounds of the cost of the transformation: a 50% overhead on the initial execution and 25% on the replay. Considering the highly non-deterministic nature of the application, the transformation should be reasonably effective on real-life applications including sizable deterministic computations. Let us note that both proposals are weak forms of execution replay as previously described. They do not allow for a parallel replay of the execution with an equivalent scheduling.

Execution replay makes it possible to remove indeterminism from parallel tracing, allowing cyclical debugging, while limiting the probe effect. It also makes it possible to relate different kinds of traces of the same program, leading to a better integration of the debugging tools. Leu and Schiper [117] show, for instance, how to integrate visualization and symbolic debugging. Execution replay should be a basic component of any parallel environment.

4.4.2 The lack of global time

It is essential that the analysis component can move forward and backward from a given event. This is easy when considering sequential executions. Events are usually time stamped with a chronological event number and the trace is organized as a sequence of events. Such an organization is still possible with parallel executions. It requires either a global (possibly logical) clock, easily available on a shared memory machine, or totally ordered logical clocks, as defined by Lamport [111]. An interesting alternative consists of organizing the global trace

as a set of local traces connected together through links corresponding to communication between the parallel activities, in the spirit of the history tapes of execution replay. Such an organization is more accurate in that it does not impose an arbitrary total order on the trace but keeps the causal dependencies between the events. A further step would be to time stamp the events with partially ordered logical clocks [79, 157], or to recreate these time stamps at analysis time (in case of post mortem analysis of a full trace), and take advantage of the applicability of these clocks to concurrency measures and global state analysis.

4.5 Connection with the analysis component

The model given in Figure 1 is a conceptual model only. It is very costly to systematically extract and send all the pieces of possibly interesting information regarding specifications, sources, and executions to the analysis module, each time something new is occurring. Traces and sources can be very large, and any particular analysis requires, in general, only selected information. Thus, communication between the two components has to be driven by the analysis component.

Partial solutions exist for analysis of executions traces. Ducassé has designed a trace query mechanism which optimizes control flow analysis [61]. Kusalik and Oster provide parameterizations of the extraction components, based on meta-interpreters [107]. Solutions derived for imperative languages in the Dalek system [147] could be useful for logic programming.

The problem, even if less acute, also exists for source access. When program sources get really large it is inappropriate to parse them each time a property has to be derived. Database systems could be used to store already parsed pieces of code.

5 Some user-oriented program analyses

The following list of analyses focuses on aspects which have not been given much attention so far in the community: browsing, explanations, profiling, and reverse-engineering. It does not pretend to be exhaustive. In particular, we conjecture that many existing static analyses developed within the framework of abstract interpretation (see, for example, [132, 21, 39, 22, 89]) could derive properties interesting for end users.

Our aim is not to discuss how analyses are implemented but to give an indication of what they can be used for in the context of programming environments. Debugging and testing, which are particularly important for programming environments, are addressed in the next section.

5.1 Browsing

For experienced programmers it is usually easy to gain local understanding of programs. This is especially true in Prolog where variables are local to a clause and where most systems now have some sort of modularity. When the size and number of modules increase

programs become harder and harder to understand. Users need some support to gain a global understanding, e.g. of the relationships between predicates and between modules.

A strong ergonomical requirement is that, at each moment, a user should not have to face more than *seven* items of information (plus or minus two) [133]. The nature of the items depends on the skill of the concerned user.

Browsers help users traverse stepwise both source and executions. Step-by-step tracing, for example, is a very primitive way of browsing. Executions and sources can be abstracted so that at each browsing step the user sees a “chunk” of information which gives meaningful information of reasonable size.

The issue is then to define abstraction criteria. For Prolog executions some necessary abstraction criteria have been identified by Bergantz and Hassell [11]: control flow, data flow, data structure, and function relations. Control flow, data flow and data abstracts have been implemented in the Opium system [58], together with failure and endless loop analyses which produce symptom oriented abstracts. Höök et al. propose a recursion abstract and a dataflow abstract [95]. Top-down zooming is a common browsing strategy [5, 70, 123]. It starts from the top most goals and recursively zooms into their executions whenever the analysis requests it. Most graphic systems discussed in the visualization section provide browsing facilities. The transparent Prolog Machine (TPM) has two levels of granularity and provides ad hoc abstraction mechanisms [16]. In a related field, [122] restructures proofs made by theorem provers. For imperative languages, the Traceview system [128] provides a number of abstract points of view.

The usual source browser is a cross referencer as, for example, Xref, written by Dave Bowen et al., and part of the AIAI Prolog library. To our point of view, much more could be done in source browsing than cross referencing. For example, it would be very helpful, when looking at recursive predicates, to automatically abstract away all the arguments which do *not* take part in the recursion.

5.2 Explanations

Explanation systems are usually based on high-level browsing but go further in the direction of users.

Explanations are essential for teaching purposes. Novice programmers have misconceptions about the language they use [189, 94]. This can be overcome by an appropriate model [25], or by tutoring systems [43]. Explanations can also be very helpful for experienced programmers to debug programs [197, 72, 59].

However, explanations are mostly required to help end-users of sophisticated programs to accept the results of computations. The problem is well-known for expert systems but it seems to be identical for other areas, such as deductive databases or systems using constraints. Unless appropriate explanation tools are provided the new technologies will not break through.

For deductive databases, the study has started [169, 193, 172, 4], but people are mainly busy extracting the resulting proof tree. No sophisticated explanation system has been built yet.

Logic programming has been used to support explanations for expert systems [42, 73, 163, 175, 196]. Conversely, studies of explanations for expert systems could certainly contribute to the explanations of logic programs, see, for example, [159, 178, 149]

5.3 Profiling

Profiling consists of investigating where the programs consume most of their time or space in order to know, for example, what needs to be optimized by hand. Although clever compilers reduce the need for profiling, there is always a need for tools to understand where computations spend time. For example, a program written in a style as declarative as possible will not in general be sufficiently efficient. However there is no need to optimize all of it (at the risk of making the program too complicated and hard to maintain). Usually 80% of a computation is spent in 20% of the code. Identifying this 20%, and only optimizing it is a significant gain in time and maintainability.

Profilers give some insight into the time spent in units of the code [87]. This information can be seen as a sort of trace. It is, however, of a lower-level than the information usually provided by tracers. Nevertheless, this information is basic and should be extracted by the same component as the other trace. Once extracted it has to be analysed in the same flexible way by the analysis component. It is a problem that existing profilers require the users to perform most of the analysis by hand.

Most existing logic programming systems provide a profiler which gives some statistics about occurrences of events and measurements of time; see, for example, ParaGraph [2] for a multiprocessor profiler. Most profilers are dynamic only. Gorlick and Kesselman [85] and Tick [181] combine static and dynamic analysis. These systems, however, lack in general the flexibility of a generic analysis component. Cohen and Carpenter [40] developed a language to inquire about the run-time behaviour of Algol programs. This language allows more general analysis to take place.

In the context of parallel systems, profiling is also useful to estimate the inherent parallelism of an application and even give parallelisation hints [177, 75, 165, 108].

5.4 Reverse-engineering

Prolog has been successfully used for reverse-engineering Cobol programs [17, 31]. Conversely, although the number of logic programs has not yet reached the critical mass which makes reverse-engineering an absolute necessity, one could imagine that reverse-engineering tools are designed for logic programming to help structure or re-structure existing code.

6 Automated debugging and testing

Debugging and testing are central issues of programming environments. We see testing as the process of analysing a program with the intent of *detecting* errors, while debugging is the process of analysing a program to *locate and fix* the detected errors.

In this section we review strategies and tools which are aimed at automating these costly activities. The boundary between testing and debugging is not easy to define. For example, some strategies are able to detect and locate at the same time. In the following, testing and debugging are therefore usually mixed.

In a survey of automated debugging covering declarative and imperative languages [62], we have identified three main strategies for automated debugging: verification with respect to specification, checking with respect to language knowledge, and filtering with respect to the error symptom. The *verification* strategy compares the actual program with some specification of the intended program. The *checking* strategy looks for suspicious places which do not comply with some explicit knowledge of the programming language. The *filtering* strategy filters out parts of the code which cannot be responsible for the error symptom. To these three strategies we have added a section on automated generation of test data, as executing programs with appropriate test data happens to be the most commonly used testing strategy in practice.

6.1 Verification with respect to specification

In the context of logic programming, the most active trend in verification is algorithmic program debugging, based on the initial work of Shapiro [168]. It compares the execution of a program against its specification. It assumes that there is such a complete specification. As it is, in general, impossible to have a complete formal specification of a program, it is actually assumed that the user can complement the formal specification and act as an “oracle”. Algorithmic debugging, in its simplest form, runs an actual computation; at each computation step the assertions derived from this computation are presented to the user, who decides whether they hold. If the program is incorrect, and if the user can correctly answer all the questions, then an erroneous component is found. It is the part of the program whose behaviour is incorrect but whose components all exhibit a correct behaviour.

Following Shapiro’s work, a number of studies have been made for Prolog [76, 123]. Algorithmic debugging is used in tutoring systems [125]. As the number of queries may be rather large, some systems use heuristics to ask users more relevant questions first [151, 28, 27, 92, 152]. Another way to reduce the number of queries is to use partial formal specifications as partial oracles [57, 56].

Algorithmic debugging has been extended to concurrent logic programming [96, 180, 120, 74], concurrent constraint logic programming [81], deductive databases [169], and Gödel [12]. Some preliminary work has been done to automatically fix programs when complete specifications are available [48, 116]. Although this assumption seems a bit unrealistic for real programs, this work could be useful for tutoring systems. Algorithmic debugging has also been adapted to imperative languages [100, 166] and lazy functional languages [141, 145].

Algorithmic debugging is an interesting technique. The current implementations, however, lack some user support. For example, users may make mistakes while acting as oracles. There must therefore be some sort of checking of their answers to validate the diagnosis of the algorithm.

Type and mode checking are other important incarnations of the verification strategy [143, 110, 18, 153].

6.2 Checking with respect to language knowledge

The checking strategy systematically parses programs and searches for language dependent errors. The knowledge of some aspects of the language is represented by consistency rules which model how these aspects should be implemented or should behave. These rules capture well-formedness knowledge which could not be encoded in the compiler. Some of the rules may assume that programming conventions are adhered to. This technique checks program sources and program executions against the consistency rules. The parts of the program which do not conform are suspicious. The rules used to detect the errors provide some explanation of the mistake.

The larger the programs, the more likely it is that they will deviate from stereotyped programming styles even if coding standards are adhered to. Hence an important issue for this technique is to choose the sensitivity of the checking: how much should a part of a program deviate from a rule to be suspected? If the technique is too sensitive it will be oversuspicious; if it is not sensitive enough it will miss many errors.

In general the suspicious code is not necessarily erroneous. Furthermore, some errors cannot be detected by a system which only uses language knowledge. Last but not least, there is no criterion to evaluate the completeness of the set of checking criteria. Therefore, there is no limit on the number of aspects to be checked. This can result in the bad situation where none of the errors have been detected while most of the program is suspicious with respect to several criteria. In such a case the programmer would have to analyse a lot of warnings to no avail.

As already mentioned, all the analyses performed by abstract interpretation could be applied here. The most typical examples are type and mode inference, see, for example, [23, 46]. They are very often thought for optimization purposes, but they can be very useful for debugging. The same applies to optimization analyses for concurrent and parallel languages, e.g. [38].

Stereotyped bug recognition has been studied for Prolog [70, 179, 125, 72]. For concurrent languages, deadlocks have been the main type of problem investigated [97, 98].

6.3 Filtering with respect to the error symptom

Once a symptom of error has been detected, the filtering strategy assumes the correctness of the program components which cannot have produced the error symptom (or which are unlikely to produce it). This strategy does not suspect any code but is aimed at reducing the amount of code which has to be examined by users.

The most important filtering technique is *slicing*, which is a symptom-driven dependency analysis [192]. So far, it has been addressing two symptoms: “wrong valued output variable” and “wrong control sequence”. Slicing prunes out the parts of the program which the wrong value, respectively the wrong control sequence, does not depend on. What remains from

the program is called a slice. Slicing according to a wrong valued variable uses a data-flow analysis, and slicing according to a wrong control sequence uses a control-flow analysis. The slice can be computed symbolically using the program source only, or it can be computed on the particular program behaviour which has exhibited the symptom under study. A slice does not always contain the erroneous statement [105], but it is always informative.

At present, there exists no slicing algorithm for Prolog or any other logic programming language. There is enough flow analyses so that this could be done promptly, see, for example, [46]. Slicing algorithms for imperative languages can be found in [192, 155, 6]. Slicing for concurrent and parallel imperative languages has also been studied [68, 35, 51].

Heuristic filtering consists of making a priori hypotheses about parts of the program which may not have caused the detected symptom in order to restrict the search space for further investigations. This technique is only acceptable if there is a possibility of backtracking over the hypotheses when further analysis does not succeed in finding the error. Pereira and Calejo have applied some heuristic filtering to logic programs, in complement with algorithmic debugging [150, 28].

6.4 Test Cases

A restrictive definition of testing is given by [138] : *Testing is the process of executing a program with the intent of finding errors.* There are other testing strategies than simply executing the program. Verification and language checking as defined earlier are good means to detect errors. However, even powerful verification and checking tools are incomplete. Executing programs with appropriate test data cannot be avoided. Therefore support is needed to generate systematic test data.

Test cases should insure that as many errors as possible are detected. It should integrate several techniques, and should be able to tell users what has been tested so far, and what confidence on which part can be assumed.

Software testing is a flourishing field of software engineering, and logic programming is used as a basis for testing other languages [86, 129, 176]. However, work on testing methodologies applied to logic programming languages is just starting [10].

Testing methodologies and techniques are essential for the acceptance of Prolog and other logic programming programs as “real” programming languages. For example, program mutation [47] should be investigated . Program mutation consists of changing small parts of a program and seeing whether the program still behaves correctly (or incorrectly).

7 Presentation and visualization of information

Once program specifications, sources and executions have been analysed and properties derived, the results have to be presented to the user. In this section we present some ergonomical presentation and visualization concepts.

We consider both textual and graphical presentation means. Indeed, some people have poor abilities to decipher 3D diagrams. It has been shown that these people are penalized

when browsing through hypertext systems [30]. Therefore environments should provide both textual and graphical tools to accommodate all kinds of users.

In the following we distinguish two types of visualization: application independent and application dependent. We believe that the issue is really to provide the support for application dependent visualization. For example, 80% of the code written for CHIP [52] applications is typically dedicated to visualization.

7.1 Application-independent visualization

We review in this section the general visualization tools.

7.1.1 Textual presentation

A straightforward and mandatory help is to display both source and trace at the same time, emphasizing their connection. For example, source oriented tracers move pointers in the source code to show where the execution currently is [154, 156, 29]. Source code, therefore, gives some context to understand execution traces. A next step is to use hypertexts as in the thesis of Calejo [27, Chapter 6].

7.1.2 Graphical presentation

Quite a number of graphical paradigms exist. These paradigms are at the end of the chain; they do not help any automation but are instead aimed at users. Enhanced proof trees have been used in the context of sequential systems. OR-parallel systems usually rely on representations of the search tree. Concurrent logic systems favour process-oriented views. Some representations are static, giving a global view, or *signature*, of the execution. Some change over time, leading to either *histories* or *animations*. There is actually no best representation, each representation gives a different view and is useful in some cases and not so useful in others. Elaborate systems combining different representations, together with new tools based on more powerful analyses, are emerging.

Execution trees and graphs In the context of sequential systems, most graphical systems are based on proof tree representations, which give an immediate picture of the proof structure and are easy to match to the clause structure of the source code, as well as to the implementation. The problem with proof trees, however, is that the information related to backtracking is lost. Ferguson diagrams [184, 50] and Aorta diagrams of the Transparent Prolog machine [71] are enhanced proof trees which retain part of the backtracking information.

In parallel systems, the focus is more on the amount and exploitation of the available parallelism. In OR-parallel systems, the focus has therefore been on the incremental building of the search tree in WAMTRACE [53], Must [177], and ParTrace [55]. Though a proof tree (or procedure invocation tree) can be used to represent an AND-parallel execution, as in VISTA [181], a graph structure has to be used to represent the synchronisation induced by

the parallel execution of AND branches. VisAndOr [33] uses such a structure in a generic way; it has been designed in order to accommodate different basic execution models (OR-parallelism, restricted AND-parallelism, and determinate dependent AND-parallelism), as well as to explore representations for combinations of these models.

The static call graph is another means used to graphically give some context to a computation. In [112], Lazzeri first builds a static call graph and, as source oriented tracers, points to nodes of this graph as the execution proceeds.

A process-oriented view Switching from the standard left to right computation rule of Prolog to coroutines [140] introduces the idea of goals “communicating” data back and forth, and suggests visualizing a logic program as a dynamic network of processes communicating via shared variables, i.e. visualizing the so-called process reading of logic programs [186]. Such a representation has been mainly used in the context of concurrent logic languages [74, 183, 127, 41].

Histories, animations, and program signatures There are many ways to look at the above-mentioned representations. Two basic choices are between a dynamic view versus a static view and a local versus a global view. Each view has its advantages.

For instance, the *animation* of the search tree offered by OR-parallel systems is a sophisticated browsing mechanism which turns out to be very good for observing the instantaneous state of a system but is not able to capture patterns of behaviour over time.

Such global patterns can be better analysed from static *program signatures* such as the ones generated by VisAndOr and VISTA. An interesting feature of VISTA is that procedure-invocation trees of concurrent logic programs are displayed by using radial coordinates and condensation. This indeed enables to better use the space on the screen, and give global pictures of executions which are still reasonably “readable” in spite of the number of nodes displayed.

In [41], Conlon and Gregory present a set of local views on query variables, shared variables between processes (channels), and processes, which facilitates the debugging of Parlog programs. In particular, the instantiation of shared variables can be displayed either in *film* format or in *snapshot* format. The film format is an animation of the successive variable instantiations. As animation in general, it does not make it easy to follow the history of the variable over time. When this is necessary, the snapshot format, listing a history of the variable states, can be used. The same kind of idea is also applied to tracing individual processes.

A better representation for parallel and concurrent executions is a process (or agent) oriented view. The sequential parts are abstracted away and the communication parts are emphasized

Integration and abstraction More powerful environments can be built by combining the above-mentioned visualizations and better integrating them in the overall environment. For instance, the programming environment of the OR-parallel system MUSE integrates,

together with a built-in statistics and benchmarking package, animation and global view by combining Must and VisAndOr [101]. This direction is also followed by its competitor system, Aurora [126]. For the time being, visualization is performed, in both systems, off-line. A further step would be to use, as suggested above, execution replay techniques to synchronize visualization and an execution replay.

A key to a better use of visualization is also the provision of higher-level analyses. This is illustrated by the work on ParSee [108], which, in the context of performance debugging, is able to characterize a predicate by a single coloured line combining three abstract measures related to granularity and scheduling.

Graphical languages Following strict visual requirements, a graphical representation of Prolog predicates is used as a basis for a graphical language [109]. It can be animated [160].

Also based on predicate representation, a graphical programming language is underway for concurrent constraint logic programming [99]. It is used for writing, animating and debugging programs.

7.2 Application-dependent visualization

As for explanations, the main issue of visualization is actually to support the end-users of the applications. We believe that the problem is not so much to provide nice representations at the level of the programming language but to provide customization mechanisms and basic graphical packages to help programmers provide appropriate graphic representations of their applications. A survey of the topic can be found in [137].

For example, the KEATS system provides graphical explanations for knowledge based systems [54]. In [37] mechanisms are introduced to enable customization of the graphics. In [15, 107] parameterization is a general design concept.

Graphical packages that have a proper interface to the environments are essential. Two such packages for Prolog are PCE [194, 195], and the graphic facilities of YAP [113, 114].

8 General programming environment issues

In this section we discuss two general problems, namely integration mechanisms and evaluation criteria.

8.1 Integration mechanisms

A major issue in programming environments in general (not only for logic programming) is to design integration mechanisms. How can we integrate, in a user-friendly and sound way, so many tools based on so many techniques? Part of the problem is that at the moment not all the sources of data are used by the tools. As illustrated by Figure 4, (partial) specifications, source code, and (traces of) executions are not used at their full potential. The three sources of data about a program should be used in synergy if a strategy requires it. The problem

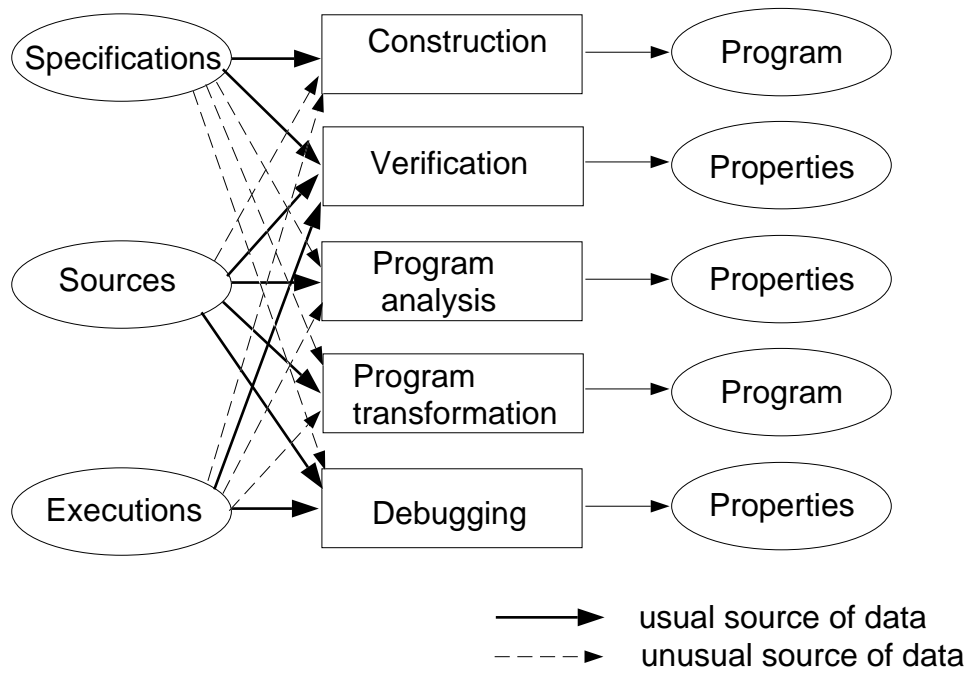


Figure 4: Current tools waste resources




Figure 5: Ideal view of programming environments

is also that tools are artificially partitioned and that they do not benefit from each other. Users want more than a list of tools, as illustrated by Figure 5.

Regarding integration, one very important requirement is flexibility. As programming environment tools are aiming at accommodating human weaknesses, they have to be flexible. No two people are the same and need exactly the same support at the same time. Some elements of solution are proposed in [60]. The analysis component provides a programming language, namely Prolog, so that static and dynamic analysis can be jointly programmed. These programs can be neatly integrated into the debugging environment thanks to an extension handler which ensures some consistency in the visibility of the extensions. In particular, on line help and insertion in the graphical interfaces are provided (almost) for free. However, much work remains to be done. The integration mechanisms of the FIELD environment for Unix should give some hints [158].

8.2 Evaluation and validation

It is difficult to draw evaluation criteria of programming environments. Since the tools are user oriented, it is not enough to measure their performance. What should be measured is their usability. The ideal way is to set up inquiries as done by Bergantz and Hassel [11] and Mulholland [136]. Such measurements are very costly. They are nevertheless very instructive. Cooperation with cognitive scientists should be reinforced to try to establish a list of criteria for the evaluation of (logic) programming environments.

Cognitive validation is a long term objective. In the meantime, a common basis should be established to test and validate the ideas at least on the same ground. In the same way that the implementation community benefited from setting up and circulating of benchmark suites, a library of important buggy and correct programs should be set up and shared among the logic programming community.

9 Conclusion

In this survey we have given a focused and maybe unusual point of view of logic programming environments. A large number of references related to logic programming environments are gathered, which show that the field, though maybe scattered, is active. We have given a framework which separates environment tools into extraction, analysis and visualization. It emphasizes that *program analysis*, both static and dynamic, is the central issue in programming environments. This point of view helps to analyse existing tools and should give some guide lines for forthcoming ones.

We have listed the achievements in logic programming, pointed out some techniques developed for other languages which should be adapted to logic programming environments, and drawn some trends for further research. Among the main achievements are algorithmic debugging, tracing for sequential Prolog, and abstract interpretation. The main techniques still missing are slicing, test case generation, and program mutation. The perspectives we see are integration, evaluation and, above all, automated static and dynamic analysis.

Programming environments are essential, however good a programming language. As program analysis is the key issue of programming environments, languages with clean semantics which enable powerful analysis have an indisputable superiority in the long term. Logic programming languages have this advantage: even Prolog with its well-known weaknesses is better than imperative languages in this area. Environments are therefore a major argument in favour of logic programming. Research on logic programming environments should therefore be mobilized as part of the marketing effort required for logic programming [185].

Acknowledgement The authors would like to thank Joachim Schimpf for his input on the “Compiled code instrumentation” section. Steven Prestwich helped on partial evaluation. Pascal Brisset, Tony Kusalik, Olivier Ridoux, Mark Wallace, and the anonymous referees gave fruitful comments on this and earlier drafts of the paper. Luke Hornof and Steven Prestwich helped with the English.

The idea of using “double quotes” to implement a simple ground representation was suggested by Lee Naish in a discussion in the newsgroup comp.lang.prolog.

The Prolog library of the AIAI at Edinburg can be retrieved by anonymous ftp (ai.ai.edinburgh.ac.uk).

This work was partly done while the authors were at the European Computer-Industry Research Centre, Munich, Germany. During this period Mireille Ducassé was partly supported by the ESPRIT Project 5291 CHIC.

References

- [1] H. Abramson and M.H. Rogers, editors. – *Meta-Programming in Logic Programming*. MIT Press, 1989. – META’88.
- [2] S. Aikawa, M. Kamiko, H. Kubo, F. Matsuzawa, and T. Chikayama. – ParaGraph: A graphical tuning tool for multiprocessor systems. – In FGCS’92 [78], pages 286–293.
- [3] ESPRIT P973 ALPES Final report, 1989. – New University of Lisbon, Portugal.
- [4] T. Arora, R. Ramakrishnan, W.G. Roth, P. Seshadri, and D. Srivastava. – Explaining program execution in deductive systems. – In S. Ceri, K. Tanaka, and S. Tsur, editors, *Proceedings of the Deductive and Object-Oriented Databases Conference*, number 760 in Lecture Notes in Computer Science. Springer-Verlag, December 1993.
- [5] E. Av-Ron. – Top-down diagnosis of Prolog programs. – Master’s thesis, Weizmann Institute of Science, Israel, 1984.
- [6] T. Ball and S. Horwitz. – Slicing programs with arbitrary control flow. – In Fritzson [80]. – AADEBUG’93.
- [7] D. Barker-Plummer. – Cliché programming in Prolog. – In Bruynooghe [20], pages 247–271. – META’90.

-
- [8] J. Barklund. – What is a meta-variable in Prolog? – In Abramson and Rogers [1], pages 383–398. – META’88.
 - [9] U. Baron, J. Chassin de Kergommeaux, M. Hailperin, M. Ratcliffe, P. Robert, J.-C. Syre, and H. Westphal. – The parallel ECRC Prolog system PEPSys: An overview and evaluation results. – In FGCS’88 [77], pages 841–850.
 - [10] F. Belli and O. Jack. – A product assurance environment for Prolog. – In Ducassé et al. [63]. – LPE’93.
 - [11] D. Bergantz and J. Hassell. – Information relationships in Prolog programs: how do programmers comprehend functionality? – *International Journal of Man-Machine Studies*, 35:313–328, 1991. – Academic Press.
 - [12] D. Binks. – Declarative debugging of abstract data types in Gödel. – In Fritzon [80]. – AADEBUG’93.
 - [13] L. Bougé, N. Choquet, L. Fribourg, and M.-C. Gaudel. – Test sets generating from algebraic specifications using logic programming. – *Journal of System and Software*, 6(4):343–360, 1986.
 - [14] A.W. Bowles and P.F. Wilk. – Tracing requirements for multi-layered meta-programming. – In Abramson and Rogers [1], pages 383–398. – META’88.
 - [15] M. Brayshaw. – An architecture for visualizing the execution of parallel logic programs. – In *Proceedings of the International Joint Conference in Artificial Intelligence*, pages 870–876. Morgan Kaufmann, 1991.
 - [16] M. Brayshaw and M. Eisenstadt. – Adding data and procedure abstraction to the Transparent Prolog Machine (TPM). – In Kowalski and Bowen [106], pages 532–547. – JICSLP’88.
 - [17] P. Breuer. – Reverse engineering in Prolog. – In N.E. Fuchs et al., editor, *Logic Programming in Action*, volume 636 of *Lecture Notes in Artificial Intelligence*, pages 290–302. Springer-Verlag, 1992.
 - [18] F. Bronsard, T.K. Lakshman, and U.S. Reddy. – A framework of directionality for proving termination of logic programs. – In K. Apt, editor, *Proceedings of the Joint International Conference and Symposium on Logic Programming*, pages 321–335, Washington D.C., USA, November 1992. MIT Press. – JICSLP’92.
 - [19] A. Bruffaerts and E. Henin. – Proof trees for negation as failure: yet another Prolog meta-interpreter. – In Kowalski and Bowen [106], pages 343–358. – JICSLP’88.
 - [20] M. Bruynooghe, editor. – *Proceedings of the second workshop on Meta-programming in Logic*, Belgium, 1990. K.U. Leuven, Department of Computer Science. – META’90.

-
- [21] M. Bruynooghe. – A practical framework for the abstract interpretation of logic programs. – *Journal of Logic Programming*, 10(2):91–124, 1991.
- [22] M. Bruynooghe. – *Special issue on Abstract Interpretation*, volume 13. – *Journal of Logic programming*, July 1992.
- [23] M. Bruynooghe and G. Janssens. – An instance of abstract interpretation integrating type and mode inference. – In Kowalski and Bowen [106], pages 669–683. – JICSLP’88.
- [24] M. Bruynooghe and M. Wirsing, editors. – *International Symposium on Programming Language Implementation and Logic Programming*, volume 631 of *Lecture Notes in Computer Science*. Springer-Verlag, August 1992. – PLILP’92.
- [25] A. Bundy, H. Pain, P. Brna, and L. Lynch. – A proposed Prolog story. – D.A.I. Research Paper 283, Department of Artificial Intelligence, University of Edinburgh, 1986.
- [26] L. Byrd. – Understanding the control flow of Prolog programs. – In Tärnlund, editor, *Logic Programming Workshop*, Debrecen, Hungary, 1980.
- [27] M.C. Calejo. – *A framework for declarative Prolog debugging*. – PhD thesis, New University of Lisbon, Portugal, March 1992.
- [28] M.C. Calejo and L.M. Pereira. – Declarative source debugging. – In *Proceedings of the 5th Portuguese AI conference*, volume 541 of *Lecture Notes in Computer Science*. Springer-Verlag, 1991.
- [29] A. Callebaut and B. Demoen. – Program sources as model for debugging in Prolog. – In Ducassé et al. [64]. – LPE’90.
- [30] F.R. Campagnoni and K. Ehrlich. – Information Retrieval using a hypertext-based help system. – *SIGIR Forum*, pages 212–220, June 1989. – Proceedings of the 12th Annual International ACM SIGIR Conference on Research and Development in Information retrieval.
- [31] G. Canfora, A. Cimitile, and U. de Carlini. – A logic-based approach to reverse engineering tools production. – *IEEE Transactions on Software Engineering*, 18(12):1053–1064, December 1992.
- [32] M. Carlsson, E. Lusk, and P. Szeredi. – Smoothing rough edges in Aurora. – In *Proceedings of the First COMPULOG-NOE Area Meeting on Parallelism and Implementation Technology*. Technical University of Madrid, Spain, May 1993.
- [33] M. Carro, L. Gomez, and M. Hermenegildo. – Some paradigms for visualizing parallel execution of logic programs. – In D.S. Warren, editor, *Proceedings of the International Conference on Logic Programming*, Budapest, Hungary, 1993. MIT Press.

-
- [34] A. Casson. – Event abstraction debuggers for layered systems in Prolog. – In *Proceedings of the UK Logic Programming Conference*. Association for Logic Programming-UK Branch, March 1990.
- [35] J. Cheng. – Slicing concurrent programs. – In Fritzson [80]. – AADEBUG'93.
- [36] C. Cherniak. – Undebuggability and cognitive science. – *Communications of the ACM*, 31(4):402–412, April 1988.
- [37] J.-L. Cochard. – A graphical representation environment for complex Prolog structures. – In Ducassé et al. [64]. – LPE'90.
- [38] M. Codish, D. Dams, and E. Shapiro. – Automatic detection of reply variables in concurrent logic programs. – In Bruynooghe [20], pages 325–338. – META'90.
- [39] C. Codognet, P. Codognet, and M. Corsini. – Abstract interpretation of concurrent logic programs. – In Debray and Hermenegildo [45], pages 215–232. – NACL'90.
- [40] J. Cohen and N. Carpenter. – A language for inquiring about the run-time behaviour of programs. – *Software-Practice and Experience*, 7:445–460, 1977.
- [41] T. Conlon and S. Gregory. – Debugging tools for concurrent logic programming. – *The Computer Journal*, 35(2):157–169, 1992. – J. Wiley and Sons.
- [42] M.J. Coombs and J. Alty. – Expert systems: an alternative paradigm. – In M. Coombs, editor, *Developments in Expert Systems*, pages 135–157. Academic Press, London, 1984.
- [43] M.J. Coombs, R.T. Hartley, and J.G. Stell. – Debugging user conceptions of interpretation processes. – In *Proceedings of the AAAI'86*, pages 303–307. Morgan Kaufmann, 1986.
- [44] J. Chassin de Kergommeaux, D. Peterson, W. Rapp, and H. Westphal. – The implementation of PEPSys on a MX-500 multiprocessor. – Technical Report CA-38, ECRC, 1988.
- [45] S.K. Debray and M. Hermenegildo, editors. – *Proceedings of the North American Conference on Logic Programming*, Austin, USA, 1990. MIT Press. – NACL'90.
- [46] S.K. Debray. – Static inference of modes and data dependencies in logic programs. – *ACM Transactions on Programming Languages and Systems*, 11(3):418–450, July 1989.
- [47] R.A. DeMillo, W.M. McCracken, R.J. Martin, and J.F. Passafiume. – *Software testing and evaluation*. – Benjamin/Cumming, Menlo Park, 1987.
- [48] N. Dershowitz and Y. Lee. – Deductive debugging. – In SLP'87 [171], pages 298–306.

-
- [49] Y. Deville and K.-K. Lau. – Logic program synthesis. – *Journal of Logic Programming*, 1994. – Also Research Report RR 93-19, Université Catholique de Louvain, Unité d’informatique.
- [50] A.D. Dewar and J.G. Cleary. – Graphical display of complex information within a Prolog debugger. – *International Journal of Man-Machine Studies*, 25(5):503–521, November 1986. – Academic Press.
- [51] C. Diehl, C. Jard, and J.-X. Rampon. – Reachability analysis on distributed executions. – In M.-C. Gaudel and J.-P. Jouannaud, editors, *Proceedings of the International Conference on Theory and Practice of Software Development (TAPSOFT’93)*, volume 668 of *Lecture Notes in Computer Science*, pages 629–643. Springer-Verlag, 1993.
- [52] M. Dinçbas, P. Van Hentenryck, H. Simonis, A. Aggoun, T. Graf, and F. Berthier. – The constraint logic programming language CHIP. – In FGCS’88 [77].
- [53] T. Disz and E. Lusk. – A graphical tool for observing the behavior of parallel logic programs. – In SLP’87 [171], pages 46–53.
- [54] J. Domingue and M. Eisenstadt. – A new metaphor for the graphical explanation of forward-chaining rule execution. – In *Proceedings of the International Joint Conference on Artificial Intelligence*, pages 129–134. Morgan Kaufmann, 1989.
- [55] M. Dorochevsky and J. Xu. – ElipSys parallel execution tracer. – Internal Report 91-7i, ECRC, February 1992.
- [56] W. Drabent, S. Nadjm-Tehrani, and J. Maluszynski. – The use of assertions in algorithmic debugging. – In FGCS’88 [77], pages 573–581.
- [57] W. Drabent, S. Nadjm-Tehrani, and J. Maluszynski. – Algorithmic debugging with assertions. – In Abramson and Rogers [1], pages 383–398. – META’88.
- [58] M. Ducassé. – Abstract views of Prolog executions in Opium. – In Saraswat and Ueda [162], pages 18–32. – ILPS’91.
- [59] M. Ducassé. – Analysis of failing Prolog executions. – In *Proceedings of the ICLP’91 Workshop on Logic Programming Environments*, Paris, June 1991. – Technical Report LIFO N 91-61, University of Orléans, France.
- [60] M. Ducassé. – *An extendable trace analyser to support automated debugging*. – PhD thesis, University of Rennes I, France, June 1992. – European Doctorate.
- [61] M. Ducassé. – A general trace query mechanism based on Prolog. – In Bruynooghe and Wirsing [24], pages 400–414. – PLILP’92.
- [62] M. Ducassé. – A pragmatic survey of automated debugging. – In Fritzon [80]. – AADEBUG’93.

-
- [63] M. Ducassé, B. Le Charlier, Y.-J. Lin, and U. Yalcinalp, editors. – *Proceedings of ILPS'93 Workshop on Logic Programming Environments*, Vancouver, October 1993. Publication IRISA, Campus de Beaulieu, 35042 Rennes, France. – LPE'93.
- [64] M. Ducassé, A.-M. Emde, T. Kusalik, and J. Levy, editors. – *Proceedings of ICLP'90 Workshop on Logic Programming Environments*, Eilat, June 1990. Technical Report, ECRC IR-LP-31-25, European Computer-Industry Research Centre. – LPE'90.
- [65] M. Ducassé and G. Ferrand, editors. – *Proceedings of ICLP'91 Workshop on Logic Programming Environments*, Paris, June 1991. Technical Report, University of Orléans, France, LIFO N 91-61. – LPE'91.
- [66] M. Ducassé, Y.-J. Lin, and L.Ü. Yalcinalp, editors. – *Proceedings of IJCSLP'92 Workshop on Logic Programming Environments*, Washington, November 1992. Technical Report TR 92-143, Case Western Reserve University, Cleveland. – LPE'92.
- [67] M. Ducassé and J. Noyé. – Tracing compiled Prolog code. – Publication IRISA- In preparation, 1994.
- [68] E. Duesterwald, R. Gupta, and M.L. Soffa. – Distributed slicing and partial re-execution for distributed programs. – In *Proc. of the Fifth Workshop on Languages and Compilers for Parallel Computing*, volume 589 of *Lecture Notes in Computer Science*. Springer-Verlag, 1992.
- [69] M. Eisenstadt. – A powerful Prolog trace package. – In *Proceedings of the 6th European Conference on Artificial Intelligence*. North-Holland, 1985, September 1984.
- [70] M. Eisenstadt. – Retrospective zooming, a knowledge based tracing and debugging methodology for logic programming. – In *Proceedings of the 9th International Joint Conference on Artificial Intelligence*. Morgan Kaufmann, 1985.
- [71] M. Eisenstadt and M. Brayshaw. – The Transparent Prolog Machine(TPM): an execution model and graphical debugger for logic programming. – *Journal of Logic Programming*, 5(4):277–342, 1988.
- [72] A.-M. Emde and M. Ducassé. – Automated debugging of non-terminating Prolog programs. – In *Proceedings of the ICLP'90 Workshop on Logic Programming Environments*, June 1990. – Technical Report, ECRC IR-LP-31-25.
- [73] A. Eriksson and A.-L. Johansson. – Neat explanation of proof trees. – In *Proceedings of the 9th International Joint Conference on Artificial Intelligence*, pages 379–381. Morgan Kaufmann, 1985.
- [74] Y. Feldman and E. Shapiro. – Temporal debugging and its visual animation. – In Saraswat and Ueda [162], pages 3–17. – ILPS'91.

-
- [75] M.J. Fernández, M. Carro, and M. Hermenegildo. – IDEal Resource Allocation (IDRA): A technique for computing accurate ideal speedups in parallel logic languages. – Technical Report FIM26.3/AI/92, Computer Science Faculty, Technical University of Madrid, September 1992.
- [76] G. Ferrand. – Error diagnosis in logic programming, an adaptation of E.Y. Shapiro’s method. – *Journal of Logic Programming*, 4:177–198, 1987.
- [77] *Proceedings of the International Conference on Fifth Generation Computer Systems*, Tokyo, Japan, December 1988. ICOT.
- [78] *Proceedings of the International Conference on Fifth Generation Computer Systems*, Tokyo, Japan, June 1992. ICOT.
- [79] C. Fidge. – Logical time in distributed computing systems. – *IEEE Computer*, pages 28–33, August 1991.
- [80] P. Fritzson, editor. – *Proceedings of the First Workshop on Automated and Algorithmic Debugging*, volume 749 of *Lecture Notes in Computer Sciences*, Linköping, Sweden, May 1993. Springer-Verlag. – AADEBUG’93.
- [81] M.P.J. Fromherz. – Declarative debugging of concurrent constraint programs. – In Fritzson [80]. – AADEBUG’93.
- [82] N.E. Fuchs. – Specifications are (preferably) executable. – *Software Engineering Journal*, pages 323–334, September 1992. – IEE Publications. Also Technical Report 91.10, Department of Computer Science, University of Zurich.
- [83] H. Gaifman, M.J. Maher, and E. Shapiro. – Replay, recovery, replication and snapshots of nondeterministic concurrent programs. – In *Proceedings of the International Conference in Principles of Distributed Computing*, Montreal, Canada, 1991. ACM Press.
- [84] J. Gallagher. – Transforming logic programs by specialising interpreters. – In *7th European Conference on Artificial Intelligence*, pages 109–122. North-Holland, 1986.
- [85] M.M. Gorlick and C.F. Kesselman. – Timing Prolog programs without clocks. – In SLP’87 [171], pages 426–432.
- [86] M.M. Gorlick, C.F. Kesselman, D.A. Marotta, and D.S. Parker. – Mockingbird: A logical methodology for testing. – *The Journal of Logic Programming*, 8(1 and 2):95–120, January / March 1990.
- [87] S.L. Graham, P.B. Kessler, and M.K. McKusick. – gprof: A call graph execution profiler. – In *Proceedings of the SIGPLAN ’82 Symposium on Compiler Construction*, pages 120–126, June 1982. – SIGPLAN Notices, Vol. 17, No. 6.

-
- [88] M. Hanus and B. Josephs. – A debugging model for functional logic programs. – In M. Bruynooghe and J. Penjam, editors, *Proceedings of the Programming Language Implementation and Logic Programming Symposium*, volume 714 of *Lecture Notes in Computer Science*. Springer-Verlag, 1993.
- [89] M.V. Hermenegildo. – Abstract interpretation and its applications. – Tutorial presented at JICSLP'92, November 1992. – Technical University of Madrid, Spain.
- [90] P.M. Hill and J.W. Lloyd. – The Gödel programming language. – Technical Report CSTR-92-27, Department of Computer Science, University of Bristol, October 1992.
- [91] P.M. Hill and J.W. Lloyd. – Analysis of meta-programs. – In Abramson and Rogers [1], pages 23–51. – META'88.
- [92] V. Hirunkitti and C.J. Hogger. – A generalised query minimisation for program debugging. – In Fritzson [80]. – AADEBUG'93.
- [93] D. Hoffman and P. Strooper. – Automated module testing in Prolog. – *IEEE Transactions on Software Engineering*, 17(9):934–943, September 1991.
- [94] K. Höök, J. Taylor, and B. du Boulay. – Redo "TRY ONCE and PASS": The influence of complexity and graphical notation on novices' understanding of Prolog. – *Instructional Science*, 19:337–360, 1990. – Kluwer Academic Publishers.
- [95] K. Höök, A. Wærn, and H. Pain. – Possible extensions to the Byrd Box tracer aimed at experts. – Poster presented at the workshop of the Psychology of Programming Interest Group, UK, January 1992.
- [96] M.M. Huntbach. – Algorithmic Parlog debugging. – In SLP'87 [171], pages 288–297.
- [97] Y. Inamura and S. Onishi. – A detection algorithm of perpetual suspension in KL1. – In Warren and Szeredi [191], pages 18–30. – ICLP'90.
- [98] S. Ishida and T. Chikayama. – Programming environment of PIMOS. – In Ducassé et al. [66], pages 48–54. – LPE'92.
- [99] K.M. Kahn. – Concurrent constraint programs to parse and animate pictures of concurrent constraint programs. – In FGCS'92 [78].
- [100] M. Kamkar, N. Shahmehri, and P. Fritzson. – Bug localization by algorithmic debugging and program slicing. – In P. Deransart and J. Maluszynski, editors, *Proceedings of the Programming Language Implementation and Logic Programming Symposium*, volume 456 of *Lecture Notes in Computer Science*, Linköping, Sweden, August 1990. Springer-Verlag.
- [101] R. Karlsson. – *A High-Performance OR-Parallel System*. – PhD thesis, The Royal Institute of Technology, Stockholm, Sweden, March 1992.

-
- [102] P. Klint, T. Reps, and G. Snelling. – Programming Environments. Report on an International Workshop at Dagstuhl Castle. – *ACM SIGPLAN Notices*, 27(11):90–96, November 1992.
- [103] J. Komorowski. – Partial evaluation as a means for inferencing data structures in an applicative language: a theory and implementation in the case of Prolog. – In *Proceedings of the Symposium on Principles of Programming Languages*, pages 255–267. ACM, 1982.
- [104] J. Komorowski, editor. – *Proceedings of the workshop on Prolog programming environments*, Sweden, 1982. University of Linköping.
- [105] B. Korel. – Identifying faulty modifications in software maintenance. – In Fritzson [80]. – AADEBUG’93.
- [106] R.A. Kowalski and K.A. Bowen, editors. – *Proceedings of the Joint International Conference and Symposium on Logic Programming*, Seattle, USA, August 1988. MIT Press. – JICSLP’88.
- [107] A.J. Kusalik and G.M. Oster. – Towards a generalized graphical interface for logic programming development. – Technical Report 93-2, University of Saskatchewan, Canada, 1993.
- [108] A.J. Kusalik and S.D. Prestwich. – Programmer-oriented visualisation of parallel logic program execution. – In Ducassé et al. [63]. – LPE’93.
- [109] D. Ladret and M. Rueher. – Vlp: A visual logic programming language. – *Journal of Visual Languages and Computing*, 2:163–188, 1991. – Academic Press.
- [110] T. K. Lakshman and U. S. Reddy. – Typed Prolog: A semantic reconstruction of the Mycroft-O’Keefe type system. – In Saraswat and Ueda [162], pages 202–220. – ILPS’91.
- [111] L. Lamport. – Time, clocks, and the ordering of events in a distributed system. – *Communications of the ACM*, 21(7):558–565, July 1978.
- [112] S.G. Lazzeri. – Vizzprol: A tool for visualizing Prolog programs. – In Ducassé and Ferrand [65]. – LPE’91.
- [113] J.P. Leal. – The Ytoolkit: a Prolog approach to a user interface. – In Ducassé et al. [64]. – LPE’90.
- [114] J.P. Leal, L. Damas, and N. Moreira. – An history based interface. – In Ducassé and Ferrand [65]. – LPE’91.
- [115] T. Leblanc and J. Mellor-Crummey. – Debugging parallel programs with Instant Replay. – *IEEE Transactions on Computers*, C-36(4):471–481, April 1987.

-
- [116] Y.-J. Lee and N. Dershowitz. – Debugging logic programs using specifications. – In Fritzson [80]. – AADEBUG'93.
- [117] E. Leu and A. Schiper. – Execution replay: a mechanism for integrating a visualization tool with. – In L. Bougé, M. Cosnard, Y. Robert, and D. Trystram, editors, *Proceedings of the 2nd Joint International Conference on Vector and Parallel Processing*, volume 634 of *Lecture Notes in Computer Science*, pages 55–66. Springer-Verlag, September 1992.
- [118] E. Leu, A. Schiper, and A. Zramdini. – Efficient execution replay technique for distributed memory architectures. – In A. Bode, editor, *Proceedings of the 2nd European Conference on Distributed Memory Computing*, volume 487 of *Lecture Notes in Computer Science*, pages 315–324. Springer-Verlag, 1991.
- [119] G. Levi and M. Martelli, editors. – *Proceedings of the International Conference on Logic Programming*, Lisbon, Portugal, June 1989. MIT Press. – ICLP'89.
- [120] Y. Lichtenstein and E. Shapiro. – Abstract algorithmic debugging. – In Kowalski and Bowen [106]. – JICSLP'88.
- [121] P. Lim and P.J. Stuckey. – Meta programming as constraint programming. – In Debray and Hermenegildo [45], pages 416–430. – NACLP'90.
- [122] C. Lingenfelder. – Structuring computer generated proofs. – In *Proceedings of the 11th International Joint Conference on Artificial Intelligence*, pages 378–383. Morgan Kaufmann, 1989.
- [123] J.W. Lloyd. – Declarative error diagnosis. – *New Generation Computing*, 5(2):133–154, 1987. – Springer-Verlag.
- [124] J.W. Lloyd. – Directions for meta-programming. – In FGCS'88 [77], pages 609–617.
- [125] C-K. Looi. – Analysing novices' programs in a Prolog intelligent teaching system. – In *Proceedings of the European Conference on Artificial Intelligence*, pages 314–319, Munich, August 1988. Pitman.
- [126] E. Lusk, S. Mudambi, R. Overbeek, and P. Szeredi. – Applications of the Aurora parallel Prolog system to computational molecular biology. – In D. Miller, editor, *Proceedings of the International Logic Programming Symposium*, pages 353–369, Vancouver, Canada, October 1993. MIT Press. – ILPS'93.
- [127] M. Maeda. – Implementing a process oriented debugger with reflection and program transformation. – In FGCS'92 [78], pages 961–968.
- [128] A. Malony, D. Hammerslag, and D. Jablonowski. – Traceview: A trace visualization tool. – *IEEE Software*, pages 19–28, September 1991.

-
- [129] B. Marre. – Toward automatic test data set selection using algebraic specifications and logic programming. – In K. Furukawa, editor, *Proceedings of the International Conference on Logic Programming*, pages 202–219, Paris, France, 1991. MIT Press.
- [130] C.E. McDowell and D.P. Helmbold. – Debugging concurrent programs. – *ACM Computing Surveys*, 21(4):593–622, December 1989.
- [131] M. Meier, A. Aggoun, D. Chan, P. Dufresne, R. Enders, D. Henry de Villeneuve, A. Herold, P. Kay, B. Perez, E. van Rossum, and J. Schimpf. – SEPIA - an extendible Prolog system. – In *Proceedings of the IFIP '89*, 1989.
- [132] C. Mellish. – Abstract interpretation of Prolog programs. – In S. Abramsky and C. Hankin, editors, *Abstract Interpretation of Declarative Languages*, chapter 8, pages 181–198. Ellis Horwood Limited, 1987.
- [133] G.A. Miller. – The magical number seven, plus or minus two: Some limits on our capacity for processing information. – *Psychological Review*, 63:81–97, 1956.
- [134] L. Monteiro and A. Porto. – Contextual logic programming. – In Levi and Martelli [119], pages 284–299. – ICLP'89.
- [135] S. Moroshita and M. Numao. – Prolog computation model BPM and its debugger PROEDIT2. – In *Proceedings of the 5th Logic Programming Conference*, pages 147–158, Tokyo, June 1986. Springer-Verlag.
- [136] P. Mulholland. – The effect of graphical and textual visualisation on the comprehension of Prolog execution by novices: an empirical analysis. – Human Cognition Research Laboratory, Open University, Milton Keynes, UK, 1993.
- [137] B.S. Murray and E. McDaid. – Visualizing and representing knowledge for the end user: a review. – *International Journal of Man-Machine Studies*, 38:23–49, 1993. – Academic Press.
- [138] G.J. Myers. – *The art of software testing*. – Business data processing. John Wiley and Sons, 1979. – ISBN 0-471-04328-1.
- [139] G. Nadathur and D.A. Miller. – An overview of λ Prolog. – In Kowalski and Bowen [106], pages 810–827. – JICSLP'88.
- [140] L. Naish. – *Negation and control in Prolog*, volume 238 of *Lecture Notes in Computer Science*. – Springer-Verlag, 1986.
- [141] L. Naish. – Declarative debugging of lazy functional programs. – In Ducassé et al. [66], pages 29–34. – LPE'92.
- [142] L. Naish. – Types and the intended meaning of logic programs. – In Pfenning [153], chapter 6, pages 189–216. – ISBN 0-262-16131-1.

-
- [143] L. Naish, P.W. Dart, and J. Zobel. – The NU-Prolog debugging environment. – In Levi and Martelli [119]. – ICLP’89.
- [144] H. Nakashima, S. Tomura, and K. Ueda. – What is a variable in Prolog? – In *Proceedings of the International Conference on Fifth Generation Computer Systems*, pages 327–332. ICOT, 1984.
- [145] H. Nilsson and P. Fritzson. – Algorithmic debugging of lazy functional languages. – In Bruynooghe and Wirsing [24], pages 385–399. – PLILP’92.
- [146] J. Noyé. – An overview of the Knowledge Crunching Machine. – In M. Abdelguerfi and S. Lavington, editors, *Emerging Trends in Database and Knowledge-base Machines*. IEEE Computer Society Press, 1994.
- [147] R.A. Olsson, R.H. Crawford, and W.W. Ho. – A dataflow approach to event-based debugging. – *Software-Practice and Experience*, 21(2):209–229, February 1991. – J. Wiley and Sons.
- [148] L. Osterweil and L.A. Clarke. – A proposed testing and analysis research initiative. – *IEEE Software*, pages 89–96, September 1992.
- [149] C.L. Paris. – Generation and explanation: Building an explanation facility for the explainable expert systems framework. – In C. Paris, W. Swartout, and W. Mann, editors, *Natural Language Generation in Artificial Intelligence and Computational Linguistics*, pages 49–81. Kluwer, Boston, 1991.
- [150] L.M. Pereira. – Rational debugging in logic programming. – In E. Shapiro, editor, *Proceedings of the International Logic Programming Conference*, volume 225 of *Lecture Notes in Computer Science*, pages 203–210, London, UK, July 1986. Springer-Verlag.
- [151] L.M. Pereira and M.C. Calejo. – A framework for Prolog debugging. – In Kowalski and Bowen [106]. – JICSLP’88.
- [152] L.M. Pereira, C.V. Damasio, and J.J. Alferes. – Debugging by diagnosing assumptions. – In Fritzson [80]. – AADEBUG’93.
- [153] F. Pfenning, editor. – *Types in Logic Programming*. – MIT Press, 1992. – ISBN 0-262-16131-1.
- [154] D. Plummer. – Coda: An extended debugger for Prolog. – Technical report AI87-54, Texas University at Austin, April 1987.
- [155] A. Podgurski and L.A. Clarke. – The implications of program dependences for software testing, debugging, and maintenance. – In R.A. Kemmerer, editor, *Proceedings of the Third Symposium on Software Testing, Analysis and Verification (TAV3)*, pages 168–178. ACM-SIGSOFT, December 1989. – Software Engineering Notes, Vol 14, No. 8.

-
- [156] T. Rajan. – Apt: A principled design of an animated view of program execution for novice programmers. – In H.-J. Bullinger and B. Shackel, editors, *Human-Computer Interaction - INTERACT'87*, pages 291–296. IFIP, September 1987.
- [157] M. Raynal. – About logical clocks for distributed systems. – *Operating Systems Review (ACM SIGOPS)*, 26(1):41–48, January 1992.
- [158] S.P. Reiss. – Connecting tools using message passing in the FIELD program development environment. – *IEEE Software*, pages 57–66, July 1990.
- [159] M.C. Rousset and B. Safar. – Negative and positive explanations in expert systems. – *Applied Artificial Intelligence*, 1, 1987. – Hemisphere Publishing Corporation.
- [160] M. Rueher. – Revisiting capabilities of graphic for logic programming. – In Ducassé and Ferrand [65]. – LPE'91.
- [161] S. Safra and E. Shapiro. – Meta interpreters for real. – In H.-J. Kugler, editor, *Information Processing*, pages 271–278, Dublin, Ireland, 1986. North-Holland.
- [162] V. Saraswat and K. Ueda, editors. – *Proceedings of the International Logic Programming Symposium*, San Diego, USA, October 1991. MIT Press. – ILPS'91.
- [163] C. Saurel. – *Contribution aux Systèmes Experts : Développement d'un Cas Concret et Etude du Problème de la Génération d'Explications Négatives*. – PhD thesis, ENSAE, Toulouse, Decembre 1987.
- [164] A. Schleiermacher and J.F.H. Winkler. – The implementation of ProTest, a Prolog debugger for a refined box model. – *Software-Practice and Experience*, 1990. – J. Wiley and Sons.
- [165] D.C. Sehr and L.V. Kalé. – Estimating the inherent parallelism in Prolog programs. – In FGCS'92 [78], pages 783–790.
- [166] N. Shahmehri. – *Generalized algorithmic debugging*. – PhD thesis, Linköping University, Department of Computer Science, Sweden, 1991. – Dissertation no 297.
- [167] E. Shapiro. – The family of concurrent logic programming languages. – *ACM Computing Surveys*, 21(3):413–510, September 1989.
- [168] E.Y. Shapiro. – *Algorithmic Program Debugging*. – MIT Press, Cambridge, MA, 1983. – ISBN 0-262-19218-7.
- [169] O. Shmueli and S. Tsur. – Logical diagnosis of LDL programs. – In Warren and Szeredi [191], pages 112–129. – ICLP'90.
- [170] W. Silverman, M. Hirsch, M. Hourii, and E. Shapiro. – The Logix system user manual. – In E. Shapiro, editor, *Concurrent Prolog: Collected Papers*. MIT Press, 1987.

-
- [171] *Proceedings of the Symposium on Logic Programming*, San Francisco, USA, September 1987. IEEE Computer Society.
- [172] G. Specht. – Generating explanation trees even for negations in deductive database systems. – In Ducassé et al. [63]. – LPE'93.
- [173] L. Sterling and R.D. Beer. – Incremental flavor-mixing of meta-interpreters for expert system construction. – In *Proceedings of the 3rd Symposium on Logic Programming*, pages 20–27, Salt Lake City, 1986.
- [174] L. Sterling and E. Shapiro. – *The Art of Prolog*. – MIT Press, Cambridge, Massachusetts, 1986. – ISBN 0-262-69105-1.
- [175] L. Sterling and L.Ü. Yalcinalp. – Explaining Prolog based expert systems using a layered meta-interpreter. – In *Proceedings of the International Joint Conference on Artificial Intelligence*. Morgan Kaufmann, 1989.
- [176] P. Strooper and D. Hoffman. – Prolog testing of C modules. – In Saraswat and Ueda [162], pages 596–610. – ILPS'91.
- [177] C. Svensson and J. Sundberg. – MUSE trace. A graphic tracer for Or-parallel Prolog. – Technical Report T90003, SICS, Sweden, 1990.
- [178] W.R. Swartout, H. Nordin, C. Paris, and S.W. Smoliar. – Toward a rapid prototyping environment for expert systems. – In *Proceedings of the 13th German Workshop on Artificial Intelligence*. Springer-Verlag, 1989.
- [179] H. Takahashi and E. Shibayama. – Preset - a debugging environment for Prolog. – In *Logic Programming Conference*, volume 221 of *Lecture Notes in Computer Science*, pages 90–99, Tokyo, Japan, 1985. Springer-Verlag.
- [180] A. Takeuchi. – Algorithmic debugging of GHC programs and its implementation in GHC. – In E. Shapiro, editor, *Concurrent Prolog: Collected Papers*, pages 180–196. MIT Press, Boston, 1987.
- [181] E. Tick. – Visualizing parallel logic programs with VISTA. – In FGCS'92 [78], pages 934–942.
- [182] G. Tobermann and C. Beckstein. – What's in a trace: The box model revisited. – In Fritzson [80]. – AADEBUG'93.
- [183] R. Trehan. – A process based tracer for KL1 on PIM. – In Ducassé and Ferrand [65]. – LPE'91.
- [184] M.H. van Emden. – An interpreting algorithm for Prolog programs. – In *Proceedings of the First International Conference on Logic Programming*, Marseille, France, September 1982.

-
- [185] M.H. van Emden. – Marketing for logic programming. – *Newsletter of the Association for Logic Programming*, 6(1):3–4, February 1993.
- [186] M.H. van Emden and G.J. de Lucena. – Logic programming. – In Clark and Tärnlund, editors, *Logic Programming*, chapter Predicate logic as a language for parallel programming, pages 189–198. Academic Press, 1982.
- [187] F. van Harmelen. – Definable naming relations in meta-level systems. – In A. Pettorossi, editor, *Meta-Programming in Logic, Third International Workshop, META-92 Uppsala, Sweden, June 1992, Proceedings*, volume 649 of *Lecture Notes in Computer Science*, pages 89–104. Springer-Verlag, 1992. – META'92.
- [188] E. van Rossum. – Implémentation d'un debugger Prolog. – Master's thesis, Facultés Universitaires Notre Dame de la Paix, Namur, Belgique, 1989.
- [189] M.W. Van Someren. – Beginners problems in learning Prolog. – Memo 54, University of Amsterdam, 1985.
- [190] A. Walker. – Prolog/Ex1, an inference engine which explains both yes and no answers. – In *proceedings of the International Joint Conference on Artificial Intelligence*, pages 526–528. Morgan Kaufmann, 1983.
- [191] D.H.D. Warren and P. Szeredi, editors. – *Proceedings of the International Conference on Logic Programming*, Jerusalem, Israël, June 1990. MIT Press. – ICLP'90.
- [192] M. Weiser. – Program slicing. – *IEEE Transactions on Software Engineering*, SE-10(4):352–357, July 1984.
- [193] C. Wieland. – Two explanation facilities for the deductive database management system DeDex. – In H. Kangassalo, editor, *Proceedings of the 9th Conference on Entity-Relationship Approach*, pages 189–203, 1990. – ETH Zurich.
- [194] J. Wielemaker and A. Anjewierden. – *PCE-4 Functional Overview*. – SWI, University of Amsterdam, Roetersstraat 15, 1018 WB Amsterdam, The Netherlands, 1992.
- [195] J. Wielemaker and A. Anjewierden. – *Programming in PCE/Prolog*. – SWI, University of Amsterdam, Roetersstraat 15, 1018 WB Amsterdam, The Netherlands, 1992.
- [196] L.Ü. Yalcinalp. – *Meta-programming for knowledge based systems in Prolog*. – PhD thesis, Case Western Reserve University, Cleveland, Ohio 44106, August 1991. – Technical Report TR 91-141.
- [197] L.Ü. Yalcinalp and L. Sterling. – An integrated interpreter for explaining Prolog's successes and failures. – In Abramson and Rogers [1]. – META'88.



Unité de recherche INRIA Lorraine, Technopôle de Nancy-Brabois, Campus scientifique,
615 rue du Jardin Botanique, BP 101, 54600 VILLERS LÈS NANCY
Unité de recherche INRIA Rennes, Irista, Campus universitaire de Beaulieu, 35042 RENNES Cedex
Unité de recherche INRIA Rhône-Alpes, 46 avenue Félix Viallet, 38031 GRENOBLE Cedex 1
Unité de recherche INRIA Rocquencourt, Domaine de Voluceau, Rocquencourt, BP 105, 78153 LE CHESNAY Cedex
Unité de recherche INRIA Sophia-Antipolis, 2004 route des Lucioles, BP 93, 06902 SOPHIA-ANTIPOLIS Cedex

Éditeur
INRIA, Domaine de Voluceau, Rocquencourt, BP 105, 78153 LE CHESNAY Cedex (France)
ISSN 0249-6399