

# The Verification Grand Challenge and Abstract Interpretation

Patrick Cousot<sup>1</sup>

École normale supérieure  
45 rue d'Ulm  
75230 Paris cedex 05  
France  
`Patrick.Cousot@ens.fr`

## 1 Introduction

Abstract Interpretation is a theory of approximation of mathematical structures, in particular those involved in the semantic models of computer systems [4,10,11]. Abstract interpretation can be applied to the systematic construction of methods and effective algorithms to approximate undecidable or very complex problems in computer science. The scope of application is rather large e.g. from type inference [5], model-checking [13], program transformation [14], watermarking [15] to context-free grammar parser generation [16].

In particular, abstract interpretation-based static analysis, which automatically infers dynamic properties of computer systems, has been very successful these last years to automatically verify complex properties of real-time, safety critical, embedded systems.

For example, ASTRÉE [1,2,3,17,18,25] can analyze mechanically and verify formally the absence of runtime errors in industrial safety-critical embedded synchronous control/command codes of several hundred thousand to one million of lines C.

We summarize the main reasons for the technical success of ASTRÉE, which provides directions for application of abstract interpretation to the Verification Grand Challenge [22,23].

## 2 The Static Analyzer ASTRÉE

### 2.1 Programs analyzed by ASTRÉE

ASTRÉE [1,2,3,17,18,25] is a static program analyzer aiming at proving the absence of Run Time Errors (RTE) in programs written in the C programming language<sup>1</sup>.

ASTRÉE analyzes structured C programs, without side effects in expressions, dynamic memory allocation and recursion. All other features of C are handled including arrays, structures, union types, pointers, pointer arithmetics, etc [31].

---

<sup>1</sup> C programs are analyzed after macro-expansion.

These restrictions encompass many synchronous, time-triggered, real-time, safety critical, embedded software programs as found in aerospace, automotive, customer electronics, defense, energy, industrial automation, medical device, rail transportation and telecommunications applications.

## 2.2 Specifications checked by ASTRÉE

ASTRÉE aims at proving that the C programming language is correctly used and that there can be no Run-Time Errors (RTE) during any execution in any environment. This covers:

- Any use of C defined by the international norm governing the C programming language (ISO/IEC 9899:1999) as having an undefined behavior (such as division by zero or out of bounds array indexing);
- Any use of C violating the implementation-specific behavior of the aspects defined by ISO/IEC 9899:1999 as being specific to an implementation of the program on a given machine (such as the size of integers and arithmetic overflow);
- Any potentially harmful or incorrect use of C violating optional user-defined programming guidelines (such as no modular arithmetic for signed integers, even though this might be the hardware choice);
- Any violation of optional, user-provided assertions (similar to assert diagnostics for example), to prove user-defined run-time properties.

## 2.3 Characteristics of ASTRÉE

ASTRÉE is sound, exhaustive, automatic, infinitary, efficient, trace-based, relational, specialized, domain-aware, parametric, modular and precise. More precisely:

- ASTRÉE is *sound* in that it always considers an over-approximation of all possible executions (for example with respect to the rounding of floating-point computations [33]). Since execution is undefined after some runtime errors and may have an unpredictable, implementation dependent effect (e.g. an array bound overflow might destroy code), ASTRÉE may have to assume that execution stops in case of definite runtime error (although in practice it may go on with an “undefined” behavior). If ASTRÉE can prove the absence of any runtime error, the program semantics is well-defined and the analysis is perfectly sound. Otherwise, the results produced by ASTRÉE describe correctly all executions before the first runtime error, if any;
- ASTRÉE is *exhaustive* and considers all possible run-time errors in all possible program executions. Hence ASTRÉE *never* omits to signal a potential run-time error, a minimal requirement for safety critical software;
- ASTRÉE is fully *automatic*, that is never needs to rely on the user’s help such as program decoration with inductive invariants. It may only happen that a few hypotheses on the range of variation of some inputs or the clock rate and maximal execution time may have to be specified in a separate configuration file to exclude e.g. impossible behaviors of the execution environment;

- ASTRÉE is *infinitary*, that is uses infinite abstractions which are provably more powerful than finite abstract models [12]. This implies that convergence acceleration techniques such as widening/narrowing [10] must be used to enforce termination of fixpoint iterations. Simultaneous widenings in several separate, independently designed, abstract domains are ensured to enforce convergence thanks to an appropriate cooperation between abstract domains [18];
- ASTRÉE always terminates and has shown to be *efficient* and to scale up to real size programs as found in the advanced industrial practice. Observed typical figures are about 1 to 2 hours of computation per hundred thousands lines (although the program size only is not an appropriate measure of the program analysis complexity);
- ASTRÉE is *trace-based*, that is abstracts sets of execution traces as opposed to invariance involving only sets of states. This refined abstraction considerably enhances the precision of the analysis, in particular for functions and procedures [26];
- ASTRÉE is *relational* that is keeps track of relations between the values of program data (variables, fields of structures, array elements, etc). Examples are the octagon abstract domain [28,29,30,32,27] or binary decision trees [3,17]. Contrary to attribute-independent abstract domains (such as intervals [9]), relational abstract domains are expensive (with a polynomial behavior with high degrees, if not exponential [19], in the number of abstract variables). To scale up, ASTRÉE uses program analysis directives (which insertion in the program can be automated by preliminary phases of the analysis) to determine which candidate packs of variables should be separately considered in relational abstractions [3,17].
- Like *general-purpose static analyzers*, ASTRÉE relies on programming language-related properties to point at potential run-time errors. Like *specialized static analyzers*, ASTRÉE puts additional restrictions on considered program (e.g. no recursion, no side-effect in expressions, no forward go to) and so can take specific program structures into account. For example function and pointer analysis involves no approximation at all, which would not be possible with dynamic memory allocation and recursion;
- Moreover, ASTRÉE is *domain-aware* and so knows facts about application domains that are indispensable to make sophisticated proofs. For example, ASTRÉE takes the logic and functional properties of control/command theory into account as implemented in embedded programs [3,20,21];
- ASTRÉE is *parametric* in that the degree of precision of the analysis can be adjusted either manually (through parameters or directives in the program text) or mechanically (by automatic insertion of the directives by preliminary analysis phases). This means that the performance rate (cost of the analysis/precision of the analysis) can be fully adapted to the needs of its end-users;
- ASTRÉE is *modular*. It is made of separate parts (so called *abstract domains*) that can be assembled and parameterized to build application specific analyzers, fully adapted to a domain of application or to end-user needs. Written

in OCaml, the modularization of ASTRÉE is made easy thanks to OCaml’s modules and functors;

- A consequence of undecidability in fully automatic static analysis is false alarms. Even a high selectivity rate of 1 false alarm over 100 operations with potential run-time errors leaves a number of doubtful cases which may be unacceptable for very large safety-critical or mission-critical software (for example, a selectivity rate of 1% yields 1000 false alarms on a program with 100 000 operations);

In contrast ASTRÉE, being modular, parametric and domain-aware can be made very *precise* and has shown to be able to produce *no false alarm* (and even *no alarm* after minor modifications for some critical programs), that is fully automated correctness proofs.

## 2.4 Program verification with ASTRÉE

The strength of ASTRÉE is that, despite fundamental undecidability limitations, it scales up and can automatically do (or has shown to be easily adaptable by specialists to do) complex proofs of absence of RTE for the considered family of synchronous control/command software. Such proofs are large, complex and subtle, even more than the program itself, whence well beyond human capacity, even using provers or proof assistants.

This strength comes from a careful, domain-specific design of the abstract interpretation. Any imprecise abstraction that would not be able to express and automatically infer, without loss of information, an *inductive* invariant which is necessary to prove absence of RTE for any program in the considered family would inexorably produce false alarms and in practice many, because of cascaded dependencies. On the other hand, an abstraction that would be too precise for the objective of proving absence of RTE in the considered family of programs would lead to excessive computational and memory costs. Essentially ASTRÉE has demonstrated in practice that for a specific program property (absence of RTE) and a specific family of programs (synchronous control/command C programs) it is possible to find an abstract interpretation of the program which encompasses all necessary inductive proofs at reasonable costs.

This strength is also the weakness of ASTRÉE. Since ASTRÉE produces “miracles” on the considered family of properties and programs, end-users would like it to produce very good results on any C program. Obviously this is impossible since the abstractions considered in ASTRÉE will miss the inductive invariants which are out of its precisely defined scope. However, new abstractions can be explored outside the current scope of ASTRÉE and easily incorporated in the static analyzer.

## 3 Directions for application of abstract interpretation to the Verification Grand Challenge

In light of the ASTRÉE, we propose a few directions for application of abstract interpretation to verification.

### 3.1 Program verification

“A program verifier uses automated mathematical and logical reasoning to check the consistency of programs with their internal and external specifications” [23]. Following E.W.D. Dijkstra, there is a clear distinction between the verification or proof of the presence of bugs (that is “testing” or “debugging”) from the verification or proof of the absence of bugs (that is “correctness verification” or “verification” for short). Of course the Verification Grand Challenge addresses the *correctness verification* only since the real challenge should be to find the *last* bug.

### 3.2 Error tracing

Nevertheless, bugs have to be considered in the development process. When an automatic verification system signals an error, it is important to be able to trace the origin of the error, in particular to determine whether it is a bug or a false alarm. Abstract slicing may be useful to trace back the part of the computation which is involved in the bug/false alarm [38,39]. Then constraint solving techniques can help finding an actual counter-example. However, finding counter-examples can be extremely difficult, if not impossible, e.g. when tracking the consequences of accumulating rounding errors after hours of floating point computations.

### 3.3 Program semantics

A program is checked with respect to a semantics that is a formal description of its computations. Numerous semantics have been proposed which differ in the level of abstraction at which they describe computations (e.g. sets of reachable states versus computation histories) and in the method for associating computations to programs (e.g. by induction on an abstract syntax using fixpoints versus using rule-based formal systems). These semantics can be organized in a hierarchy by abstract interpretation [7] so that different analyzers can rely on different semantics which can be formally guaranteed to be coherent, at various levels of abstractions.

In practice, although norms do exist for programming languages like C, they are of little help because too many program behaviors are left unspecified. So one must rely on compilers, linkers, loaders and machines to know, e.g. the exact effect of evaluating an arithmetic expressions. Since the Verification Grand Challenge addresses “significant software products”, it is clear that methods for defining the semantics of programs are needed, at a level of precision which is compatible with the implementation.

An approach could be, like in ASTRÉE for absence of side effects in expressions, to reject programs for which this compatibility cannot be formally guaranteed. The abstraction methods to do so, might then be part of the programming language semantics, the restricted language being a subset of a larger existing language where undefined behaviors have been excluded. Such programming

norms limiting the use of obscure, error-prone and/or non-portable features of programming languages tend to be more widely accepted in practice.

### 3.4 Compilation

That the semantics of a programming language precisely reflects program executions depends upon the correctness of the translation into machine code. So that verification of the compiler [24], of the object code [35,36] or of the translation [37] must be part of the complete verification process.

### 3.5 Specification

The program semantics restricts the verification to properties that can be expressed in terms of this semantics. The specifications (such as invariance, safety, security, liveness) further restricts the verification process to specific properties. Specifications themselves translate external requirements in terms of program computations. Thus it is necessary to define adequate specification languages, their semantics with respect to that of the programming language. This ranges from implicit internal specifications (like absence of runtime errors as defined by the semantics of the programming language) to arbitrary complex specification languages.

Specifications cannot be simply be considered as correct, since in practice they are not or e.g. only one side of interfaces satisfies the given specifications. Abstract interpretation techniques could be used both to analyze specifications and to check programs for resistance to specification unsatisfaction.

Finally, the analysis of the specification, should be checked to remain valid in the implementation, e.g. by reanalysis of the program or by translation validation.

### 3.6 Specification and verification of complex systems

More generally, specifications refer, especially in the case of embedded systems, to an external world which should be taken into account to prove the correctness of a whole system, not only the program component. Progress has to be made on the abstraction of this external, often physical world, to be compatible with the program interfaces. We envision that abstraction can be applied to the full system (program + reactive environment) although the descriptions of the program and physical part of the system are a quite different nature (e.g. continuous versus discrete). A unification of abstraction in computer science and engineering sciences must be considered to achieve the goal of full system verification [8].

### 3.7 Verification of program families

The considered programs to be verified may range from one program (with a finite specific abstraction), to a family of programs with specific characteristics, to a programming language or even a family of programming languages. A

broad spectrum verifier is likely to have many customers but also to produce too many false alarms, a recurrent complain of end-users of static analyzers. A finite abstraction can always be found for a given program and specification but discovering this abstraction amounts to making the proof [6], i.e. iteratively computing the weakest inductive argument. To get no false alarm, the consideration of families of programs for which generic, precise and efficient abstractions can be found might be a useful alternative, as was the case in ASTRÉE.

### 3.8 Required precision of verifiers

Automatic program verification requires the discovery of inductive arguments (for loops, recursion, etc). Proceeding by direct reference to the program semantics (as in refinement-based methods) amounts to the computation of the program semantics restricted to the program specification, which is not a finitary process. Abstraction is therefore necessary but leads to false alarms. The condition for absence of false alarm is that the weakest inductive argument suitable for the proof be expressible without loss of precision in the abstract (including for its transformers in the induction step) [6]. There is obviously no hope to find an abstract domain containing all of such inductive arguments, since this will ultimately amount to include e.g. all first-order predicates with arithmetic and one is back to undecidability.

### 3.9 Abstract assertions

The choice of the form of the abstract assertions depends on the considered family of programs, the nature of the considered specifications and the corresponding necessary inductive arguments. Universal representations (as terms or specific encodings of sets of states), to be used in all circumstances, are likely to be very inefficient. The specific abstract assertions are implemented as abstract domains in ASTRÉE using specific encoding and computer representations that lead to efficient manipulation algorithms. The study of efficient implementation of abstract assertions and efficient algorithms in abstract domains can certainly make significant progress, in particular by considering the domains of applications of programs.

### 3.10 Application-aware verifiers

ASTRÉE is a program verifier with a very precise scope of application that is of synchronous, real-time control command systems. It can therefore incorporate knowledge about such programs, looking e.g. for ellipsoidal assertions when encountering digital filters [20]. In absence of such domain specific knowledge, a verifier might have to look for costly nonlinear invariants.

Among the application domains that have been largely neglected by the verification community are the numerical applications involving intensive floating point computations. To be sound ASTRÉE must perform a rigorous analysis of floating point computations [30]. Further abstractions of this complex semantics are needed.

### 3.11 Abstract solvers

ASTRÉE uses sophisticated iteration techniques to propagate assertions and perform inductive steps by widening in solvers (see e.g. trace partitioning [26]). A lot of progress can still be done on abstract solvers, in particular for generic, parametric, modular and parallel [34] ones.

### 3.12 Combination of abstractions

A verification in ASTRÉE is done by parts, each part corresponding to a separate abstract domain handling specific abstract assertions, with an interaction between the parts, formalized by the reduced product [11]. So a specific version of ASTRÉE is built by incorporating a choice of abstract domains, which can be program specific, and of the corresponding interactions [18].

### 3.13 Modular analyzers

The modular design of ASTRÉE might be a useful approach to the necessity to have specific analyzers adapted to domains of applications and the need for general tools for program verification. One can imagine a large collections of abstract domains and solvers that can be combined on demand to adjust the cost/precision ratio, depending upon the proposed application of the verifier.

### 3.14 User interface

Static analyzers like ASTRÉE yield extremely complex informations on program executions that the end-user may want to understand. For that purpose, a user-interface is needed to present internal information in understandable form, at different levels of abstractions.

### 3.15 Verifier infrastructure

Once a programming, a specification and a user-interface language have been chosen, and their semantics defined, a static analyzer has to provide computer representations of these languages for the purpose of the program analysis, the specification checking and the report of the results to the end-user. Despite the difficulty to design such a general-purpose infrastructure, it should be sharable between different verifiers to accelerate experimentations and developments.

### 3.16 The verified verifier

A recurrent question about ASTRÉE is whether it has been verified and this question is likely to appear for any verifier. A verification has three phases,

1. the computation of an inductive assertion implied by the semantics and the specification which involves resolution of fixpoint inequations,



2. the verification that the assertion is indeed inductive, and
3. finally the proof that the inductive assertion implies the specification.

All phases are formally specified by abstract interpretation theory. The first phase is indeed the more complex but, from a strict soundness point of view, it does not need to be formally verified. Only the second and third phases of the verifier must be verified, which is simpler. Preliminary work on ASTRÉE shows that this is indeed possible. The verified verifier is indeed part of the Verification Grand Challenge.

### 3.17 Acceptance and dissemination of static analysis

The dissemination and widespread adoption of formal methods is confronted with economic payoff criteria. Not doing any correctness proof is, at first sight, easier and less expensive.

Regulation might be necessary to enforce the adoption of formal methods to produce safer software (e.g. in industrial norms). Static analysis, which has shown to scale up in an industrial context, is a very good candidate.

End-users might also be willing to enforce their right for verified software products. The ability to perform automatically static analyzes showing that products are not state of the art might even be a decisive argument to change present-day permissive laws regarding software reliability.

## 4 Conclusion

Abstraction, as formalized by Abstract Interpretation, is certainly central in the Verification Grand Challenge, as shown by its recent applications, that do scale up for real-life safety critical industrial applications. A Grand Challenge for abstract interpretation is to extend its scope to complex systems, from design to implementation.

## References

1. B. Blanchet, P. Cousot, R. Cousot, J. Feret, L. Mauborgne, A. Miné, D. Monniaux, and X. Rival. The ASTRÉE Static Analyzer. <http://www.astree.ens.fr/>.
2. B. Blanchet, P. Cousot, R. Cousot, J. Feret, L. Mauborgne, A. Miné, D. Monniaux, and X. Rival. Design and implementation of a special-purpose static program analyzer for safety-critical real-time embedded software, invited chapter. In T. Mogenssen, D.A. Schmidt, and I.H. Sudborough, editors, *The Essence of Computation: Complexity, Analysis, Transformation. Essays Dedicated to Neil D. Jones*, Lecture Notes in Computer Science 2566, pages 85–108. Springer, Berlin, Germany, 2002.
3. B. Blanchet, P. Cousot, R. Cousot, J. Feret, L. Mauborgne, A. Miné, D. Monniaux, and X. Rival. A static analyzer for large safety-critical software. In *Proceedings of the ACM SIGPLAN '2003 Conference on Programming Language Design and Implementation (PLDI)*, pages 196–207, San Diego, California, United States, 7–14 June 2003. ACM Press, New York, New York, United States.

4. P. Cousot. *Méthodes itératives de construction et d'approximation de points fixes d'opérateurs monotones sur un treillis, analyse sémantique de programmes (in French)*. Thèse d'État ès sciences mathématiques, Université scientifique et médicale de Grenoble, Grenoble, France, 21 March 1978.
5. P. Cousot. Types as abstract interpretations, invited paper. In *Conference Record of the Twentyfourth Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*, pages 316–331, Paris, France, January 1997. ACM Press, New York, New York, United States.
6. P. Cousot. Partial completeness of abstract fixpoint checking, invited paper. In B.Y. Choueiry and T. Walsh, editors, *Proceedings of the Fourth International Symposium on Abstraction, Reformulation and Approximation, SARA '2000*, Horseshoe Bay, Texas, United States, Lecture Notes in Artificial Intelligence 1864, pages 1–25. Springer, Berlin, Germany, 26–29 July 2000.
7. P. Cousot. Constructive design of a hierarchy of semantics of a transition system by abstract interpretation. *Theoretical Computer Science*, 277(1–2):47–103, 2002.
8. P. Cousot. Integrating physical systems in the static analysis of embedded control software, invited paper. In *Proceedings of the Third Asian Symposium on Programming Languages and Systems, APLAS '2005*, pages 135–138, Tsukuba, Japan, 3–5 November 2005. Lecture Notes in Computer Science 3780, Springer, Berlin, Germany.
9. P. Cousot and R. Cousot. Static determination of dynamic properties of programs. In *Proceedings of the Second International Symposium on Programming*, pages 106–130, Paris, France, 1976. Dunod, Paris, France.
10. P. Cousot and R. Cousot. Abstract interpretation: a unified lattice model for static analysis of programs by construction or approximation of fixpoints. In *Conference Record of the Fourth Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*, pages 238–252, Los Angeles, California, 1977. ACM Press, New York, New York, United States.
11. P. Cousot and R. Cousot. Systematic design of program analysis frameworks. In *Conference Record of the Sixth Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*, pages 269–282, San Antonio, Texas, 1979. ACM Press, New York, New York, United States.
12. P. Cousot and R. Cousot. Comparing the Galois connection and widening/narrowing approaches to abstract interpretation, invited paper. In M. Bruynooghe and M. Wirsing, editors, *Proceedings of the Fourth International Symposium Programming Language Implementation and Logic Programming, PLILP '92*, Leuven, Belgium, 26–28 August 1992, Lecture Notes in Computer Science 631, pages 269–295. Springer, Berlin, Germany, 1992.
13. P. Cousot and R. Cousot. Temporal abstract interpretation. In *Conference Record of the Twentyseventh Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*, pages 12–25, Boston, Massachusetts, United States, January 2000. ACM Press, New York, New York, United States.
14. P. Cousot and R. Cousot. Systematic design of program transformation frameworks by abstract interpretation. In *Conference Record of the Twentyninth Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*, pages 178–190, Portland, Oregon, United States, January 2002. ACM Press, New York, New York, United States.
15. P. Cousot and R. Cousot. An abstract interpretation-based framework for software watermarking. In *Conference Record of the Thirtyfirst Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*, pages 173–185,

- Venice, Italy, 14–16 January 2004. ACM Press, New York, New York, United States.
16. P. Cousot and R. Cousot. Grammar analysis and parsing by abstract interpretation, invited chapter. In T. Reps, M. Sagiv, and J. Bauer, editors, *Program Analysis and Compilation, Theory and Practice: Essays dedicated to Reinhard Wilhelm*, Lecture Notes in Computer Science 4444, pages 178–203. Springer, Berlin, Germany, 2006.
  17. P. Cousot, R. Cousot, J. Feret, L. Mauborgne, A. Miné, D. Monniaux, and X. Rival. The ASTRÉE analyser. In M. Sagiv, editor, *Proceedings of the Fourteenth European Symposium on Programming Languages and Systems, ESOP '2005, Edinburgh, Scotland*, volume 3444 of *Lecture Notes in Computer Science*, pages 21–30. Springer, Berlin, Germany, 2–10 April 2005.
  18. P. Cousot, R. Cousot, J. Feret, L. Mauborgne, A. Miné, D. Monniaux, and X. Rival. Combination of abstractions in the ASTRÉE static analyzer, invited paper. In M. Okada and I. Satoh, editors, *Eleventh Annual Asian Computing Science Conference, ASIAN 06*, Tokyo, Japan, 6–8 December 2006. Lecture Notes in Computer Science, Springer, Berlin, Germany. To appear.
  19. P. Cousot and N. Halbwachs. Automatic discovery of linear restraints among variables of a program. In *Conference Record of the Fifth Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*, pages 84–97, Tucson, Arizona, 1978. ACM Press, New York, New York, United States.
  20. J. Feret. Static analysis of digital filters. In D. Schmidt, editor, *Proceedings of the Thirteenth European Symposium on Programming Languages and Systems, ESOP '2004, Barcelona, Spain*, volume 2986 of *Lecture Notes in Computer Science*, pages 33–48. Springer, Berlin, Germany, March 27 – April 4, 2004.
  21. J. Feret. The arithmetic-geometric progression abstract domain. In R. Cousot, editor, *Proceedings of the Sixth International Conference on Verification, Model Checking and Abstract Interpretation (VMCAI 2005)*, pages 42–58, Paris, France, 17–19 January 2005. Lecture Notes in Computer Science 3385, Springer, Berlin, Germany.
  22. C.A.R. Hoare. The verifying compiler, a grand challenge for computing research. *Journal of the Association for Computing Machinery*, 50(1):63–69, 2003.
  23. C.A.R. Hoare. The verifying compiler, a grand challenge for computing research. In R. Cousot, editor, *Proceedings of the Sixth International Conference on Verification, Model Checking and Abstract Interpretation (VMCAI 2005)*, pages 78–78, Paris, France, 7–19 January 2005. Lecture Notes in Computer Science 3385, Springer, Berlin, Germany.
  24. X. Leroy. Formal certification of a compiler back-end or: programming a compiler with a proof assistant. In *Conference Record of the Thirtythird Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*, pages 42–54. ACM Press, New York, New York, United States, Charleston, South Carolina, 2006.
  25. L. Mauborgne. ASTRÉE: Verification of absence of run-time error. In P. Jacquart, editor, *Building the Information Society*, chapter 4, pages 385–392. Kluwer Academic Publishers, Dordrecht, The Netherlands, 2004.
  26. L. Mauborgne and X. Rival. Trace partitioning in abstract interpretation based static analyzer. In M. Sagiv, editor, *Proceedings of the Fourteenth European Symposium on Programming Languages and Systems, ESOP '2005, Edinburgh, Scotland*, volume 3444 of *Lecture Notes in Computer Science*, pages 5–20. Springer, Berlin, Germany, April 2–10, 2005.

27. A. Miné. The Octagon abstract domain library. <http://www.di.ens.fr/~mine/oct/>.
28. A. Miné. A new numerical abstract domain based on difference-bound matrices. In O. Danvy and A. Filinski, editors, *Proceedings of the Second Symposium PADO '2001, Programs as Data Objects*, Århus, Denmark, 21–23 May 2001, Lecture Notes in Computer Science 2053, pages 155–172. Springer, Berlin, Germany, 2001.
29. A. Miné. A few graph-based relational numerical abstract domains. In M. Hermenegildo and G. Puebla, editors, *Proceedings of the Ninth International Symposium on Static Analysis, SAS '02*, volume 2477 of *Lecture Notes in Computer Science*, pages 117–132. Springer, Berlin, Germany, 2002.
30. A. Miné. Relational abstract domains for the detection of floating-point run-time errors. In D. Schmidt, editor, *Proceedings of the Thirteenth European Symposium on Programming Languages and Systems, ESOP '2004, Barcelona, Spain*, volume 2986 of *Lecture Notes in Computer Science*, pages 3–17. Springer, Berlin, Germany, March 27 – April 4, 2004.
31. A. Miné. Field-sensitive value analysis of embedded C programs with union types and pointer arithmetics. In *Proceedings of the ACM SIGPLAN/SIGBED Conference on Languages, Compilers, and Tools for Embedded Systems, LCTES '2006*, pages 54–63. ACM Press, New York, New York, United States, June 2006.
32. A. Miné. The octagon abstract domain. *Higher-Order and Symbolic Computation*, 19:31–100, 2006.
33. A. Miné. Symbolic methods to enhance the precision of numerical abstract domains. In E.A. Emerson and K.S. Namjoshi, editors, *Proceedings of the Seventh International Conference on Verification, Model Checking and Abstract Interpretation (VMCAI 2006)*, pages 348–363, Charleston, South Carolina, United States, 8–10, January 2006. Lecture Notes in Computer Science 3855, Springer, Berlin, Germany.
34. D. Monniaux. The parallel implementation of the ASTRÉE static analyzer. In *Proceedings of the Third Asian Symposium on Programming Languages and Systems, APLAS '2005*, pages 86–96, Tsukuba, Japan, 3–5 November 2005. Lecture Notes in Computer Science 3780, Springer, Berlin, Germany.
35. X. Rival. Abstract interpretation based certification of assembly code. In L. Zuck, P. Attie, A. Cortesi, and S. Mukhopadhyay, editors, *Proceedings of the Fourth International Conference on Verification, Model Checking and Abstract Interpretation (VMCAI 2003)*, pages 41–55, , New York, New York, United States, 9–11 January 2003. Lecture Notes in Computer Science 2575, Springer, Berlin, Germany.
36. X. Rival. Invariant translation-based certification of assembly code. *International Journal on Software and Tools for Technology Transfer*, 6(1):15–37, July 2004.
37. X. Rival. Symbolic transfer functions-based approaches to certified compilation. In *Conference Record of the Thirtyfirst Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*, pages 1–13, Venice, Italy, 2004. ACM Press, New York, New York, United States.
38. X. Rival. Abstract dependences for alarm diagnosis. In *Proceedings of the Third Asian Symposium on Programming Languages and Systems, APLAS '2005*, pages 347–363, Tsukuba, Japan, 3–5 November 2005. Lecture Notes in Computer Science 3780, Springer, Berlin, Germany.
39. X. Rival. Understanding the origin of alarms in ASTRÉE. In C. Hankin and I. Siveroni, editors, *Proceedings of the Twelfth International Symposium on Static Analysis, SAS '05*, pages 303–319, London, United Kingdom, Lecture Notes in Computer Science 3672, 7–9 september 2005.