

# Reducing the Overhead of Assertion Run-time Checks via Static Analysis<sup>\*</sup>

Nataliia Stulova<sup>1</sup>

José F. Morales<sup>1</sup>

Manuel V. Hermenegildo<sup>1,2</sup>

<sup>1</sup>IMDEA Software Institute  
{nataliia.stulova,josef.morales,manuel.hermenegildo}  
@imdea.org

<sup>2</sup>School of Computer Science  
Technical University of Madrid (UPM)  
manuel.hermenegildo@upm.es

## Abstract

In order to aid in the process of detecting incorrect program behaviors, a number of approaches have been proposed which include a combination of language-level constructs (such as procedure-level assertions/contracts, program-point assertions, gradual types, etc.) and associated tools (such as code analyzers and run-time verification frameworks). However, it is often the case that these constructs and tools are not used to their full extent in practice due to a number of limitations such as excessive run-time overhead and/or limited expressiveness. Verification frameworks that combine static and dynamic techniques offer the potential to bridge this gap. In this paper we explore the effectiveness of abstract interpretation in detecting parts of program specifications that can be statically simplified to true or false, as well as the impact of such analysis in reducing the cost of the run-time checks required for the remaining parts of these specifications. Starting with a semantics for programs with assertion checking, and for assertion simplification based on static analysis information, we propose and study a number of practical assertion checking modes, each of which represents a trade-off between code annotation depth, execution time slowdown, and program safety. We also propose techniques for taking advantage of the run-time checking semantics to improve the precision of the analysis. Finally, we study experimentally the performance of these techniques. Our experiments illustrate the benefits and costs of each of the assertion checking modes proposed as well as the benefit of analysis for these scenarios.

**Keywords** Abstract Interpretation, Assertions, Run-time Checking, Verification, Logic Programming, Horn Clauses.

<sup>\*</sup>This research has been partially funded by EU FP7 agreement 318337 ENTRA, Spanish MINECO project TIN2015-67522-C3-1-R *TRACES*, and Madrid Region program M141047003 *N-GREENS*. We would also like to thank the anonymous reviewers for providing valuable comments and suggestions.

## 1. Introduction

Detecting incorrect program behaviors is an important part of the software development life cycle. It is also a complex and tedious one, in which dynamic languages bring special challenges. A number of techniques have been proposed to aid in the process, among which we center our attention on the use of language-level constructs to describe expected program behavior, and of associated tools to compare actual program behavior against expectations, such as static code analyzers/verifiers and run-time verification frameworks. Approaches that fall into this category are the assertion-based frameworks used in (Constraint) Logic Programming (Drabent et al. 1988; Puebla et al. 1997; Bueno et al. 1997; Boye et al. 1997; Hermenegildo et al. 1999; Puebla et al. 2000b; Lai 2000; Hermenegildo et al. 2005; Mera et al. 2009), soft/gradual typing approaches in functional programming (Cartwright and Fagan 1991; Findler and Felleisen 2002; Tobin-Hochstadt and Felleisen 2008; Dimoulas and Felleisen 2011; Rastogi et al. 2015; Takikawa et al. 2015, 2016), and contract-based extensions in object-oriented programming (Lampert and Paulson 1999; Leavens et al. 2007; Fähndrich and Logozzo 2011). These tools are aimed at detecting violations of the expected behavior or certifying the absence of any such violations, and often involve a certain degree of run-time testing, specially for non-trivial properties.

A practical limitation of many of these tools is that they can incur significant run-time performance overhead, even in the simple case of performing just type checks between typed and untyped parts of programs (Rastogi et al. 2015; Takikawa et al. 2016). In (Mera et al. 2009) reductions were obtained by limiting the points at which the tests are performed and the instrumentation, and by inlining, but some types of tests still incurred significant costs. Some approaches have opted for limiting the expressiveness of the assertion language in order to reduce the overhead (see (Schiller et al. 2014) for some recent case studies). Recently, some proposals have been made for reducing the run-time overhead of assertion checking based on optimizing the run-time checking mechanisms themselves, at the expense of increased memory consumption (Koukoutos and Kuncak 2014; Stulova et al. 2015): repeated checks on immutable recursive data structures are converted from execution time overhead to increased memory use via caching and/or tabling techniques.

However, despite these advances, run-time overhead often remains impractically high, specially for complex properties, such as, for example, deep data structure tests. This reduces the attractiveness of run-time checking to programmers, which may allow sporadic checking of very simple conditions, but tend to turn off run-time checking for more complex properties.

Motivated by this problem, assertion-based frameworks have been proposed where static analysis is used to minimize the number and cost of the run-time checks that need to be placed in the program to detect incorrect program behaviors. This idea was pioneered by the Ciao system (Puebla et al. 1997; Bueno et al. 1997; Hermenegildo et al. 1999; Puebla et al. 2000b; Hermenegildo et al. 2005; Pietrzak et al. 2006; Hermenegildo et al. 2012) where a number of (abstract interpretation-based) static analyses are combined in order to verify assertions to the largest extent possible at compile time, and for simplifying and reducing the number of remaining properties that that need to be introduced in the program as run-time checks. Intuitively, this model can offer a more appealing trade-off of performance vs. safety guarantees. However, while there has been evidence from use, there has been little systematic experimental work presented to date verifying this hypothesis, i.e., measuring the actual impact of analysis on reducing run-time checking overhead. For example, in (Mera et al. 2009, 2011) the overhead of run-time checking was studied but without taking into account analysis information.

In order to bridge this gap, in this work we explore the effectiveness of abstract interpretation-based compile-time analysis in detecting parts of program specifications that can be simplified before they are turned into run-time checks. Again, the objective of such simplification is to achieve a system that can detect the same (or a larger) set of incorrect behaviors in a program, but with a significant reduction in the impact on the running time of the program.

Starting with a semantics for programs with assertion checking and for assertion simplification based on analysis information obtained via abstract interpretation, we propose and study a number of practical *assertion checking modes*, each of which represents a trade-off between code annotation depth, execution time slowdown, and program behavior safety guarantees. The proposed modes are specially tailored to the scenario of annotating and pre-processing libraries to ensure their correctness prior to their use by client programs. We also define a transformation-based approach in order to implement each one of these modes.

We then concentrate on the reduction of the number of run-time tests via (abstract interpretation-based) program analysis. To this end we propose a technique that enhances analysis precision by taking into account that any assertions that cannot be proved statically will be the subject of run-time testing. We then report on an implementation of the proposed techniques (within the CiaoPP system) and study their impact in practice, by measuring the reduction in run-time checking overhead achieved.

We develop the discussion in the context of (*Horn Clause*) *Logic Programs*, to take advantage of the availability of mature program analysis and transformation tools, and a well developed assertion language and assertion processing framework (in particular, that of the Ciao system). However, we believe the results are applicable to other programming paradigms, either directly (to, e.g., other forms of declarative programming), or, following recent work, to imperative programs, via transformation into Horn Clauses (Méndez-Lojo et al. 2007; Grebenshchikov et al. 2012; Bjørner et al. 2014; Gurfinkel et al. 2015). Examples of this include cost analysis of Java bytecode programs (Navas et al. 2008, 2009), or inferring energy bounds in binaries from C-style programs (Liqat et al. 2014, 2016).

The rest of the paper is structured as follows: Section 2 presents the run-time checking part of our approach. After introducing some notation and the basic semantics in Section 2.1, Section 2.2 presents the assertion language and Section 2.3 the operational semantics with run-time checking of such assertions. Section 3 then presents the run-time assertion checking modes proposed, including a discussion of the transformations required to implement the different modes. Section 4 then addresses the issue of optimizing run-time

checks via static analysis. Section 4.1 presents the basic abstract interpretation-based analysis approach used and the memo table representation of the analysis results. Section 4.2 describes how run-time tests are optimized using the information in the analysis memo table. Section 5 then presents our approach for taking advantage of the run-time checking semantics to improve the precision of the analysis. Section 6 describes our experimental harness and presents our results for the different options (with and without analysis, with and without improved analysis precision). Section 7 finally presents our conclusions.

## 2. Run-Time Checking of Assertions

### 2.1 Basic notation and standard semantics

We revisit here some basic notation and the standard program semantics, where we use the formalization of (Stulova et al. 2015, 2014; Puebla et al. 2000b).

An *atom* has the form  $p(t_1, \dots, t_n)$  where  $p$  is a predicate symbol of arity  $n$  and  $t_1, \dots, t_n$  are terms. A *constraint* is a conjunction of expressions built from predefined predicates (such as term equations or inequalities over the reals) whose arguments are constructed using predefined functions (such as real addition). A *literal* is either an atom or a constraint. A *goal* is a finite sequence of literals. A *rule* is of the form  $H :- B$  where  $H$ , the *Head*, is an atom and  $B$ , the *body*, is a possibly empty finite sequence of literals. A *constraint logic program*, or *program*, is a finite set of rules.

The *definition* of an atom  $A$  in a program,  $\text{defn}(A)$ , is the set of variable renamings of the program rules s.t. each renaming has  $A$  as a Head and has distinct new local variables. We assume that all rule Heads are *normalized*, i.e.,  $H$  is of the form  $p(X_1, \dots, X_n)$  where the  $X_1, \dots, X_n$  are distinct free variables. Let  $\exists_L \theta$  be the constraint  $\theta$  restricted to the variables of the syntactic object  $L$ . We denote *constraint entailment* by  $\models$ , so that  $\theta_1 \models \theta_2$  denotes that  $\theta_1$  entails  $\theta_2$ . Then, we say that  $\theta_2$  is *weaker* than  $\theta_1$ .

The operational semantics of a program is given in terms of its *derivations*, which are sequences of *reductions* between *states*. A *state*  $\langle G \mid \theta \rangle$  consists of a goal  $G$  and a constraint store (or *store* for short)  $\theta$ . We use  $::$  to denote concatenation of sequences and we assume for simplicity that the underlying constraint solver is complete. A state  $S = \langle L :: G \mid \theta \rangle$  where  $L$  is a literal can be *reduced* to a state  $S'$  as follows:

1.  $\langle L :: G \mid \theta \rangle \rightsquigarrow \langle G \mid \theta \wedge L \rangle$  if  $L$  is a constraint and  $\theta \wedge L$  is satisfiable.
2.  $\langle L :: G \mid \theta \rangle \rightsquigarrow \langle B :: G \mid \theta \rangle$  if  $L$  is an atom of the form  $p(t_1, \dots, t_n)$ , for some rule  $(L :- B) \in \text{defn}(L)$ .

We use  $S \rightsquigarrow S'$  to indicate that a reduction can be applied to state  $S$  to obtain state  $S'$ . Also,  $S \rightsquigarrow^* S'$  indicates that there is a sequence of reduction steps from state  $S$  to state  $S'$ . We denote by  $D_{[i]}$  the  $i$ -th state of the derivation. As a shorthand, given a non-empty derivation  $D$ ,  $D_{[-1]}$  denotes the last state. A *query* is a pair  $(L, \theta)$ , where  $L$  is a literal and  $\theta$  a store, for which the constraint logic programming system starts a computation from state  $\langle L \mid \theta \rangle$ . The set of all derivations from the query  $Q$  is denoted  $\text{derivs}(Q)$ . A finished derivation from a query  $(L, \theta)$  is *successful* if the last state is of the form  $\langle \square \mid \theta' \rangle$ , where  $\square$  denotes the empty goal sequence. In that case, the constraint  $\exists_L \theta'$  is an *answer* to  $S$ . We denote by  $\text{answers}(Q)$  the set of answers to a query  $Q$ .

### 2.2 Assertion Language

We assume that program specifications are provided by means of assertions: linguistic constructions that allow expressing properties

of programs. In particular, we would like to specify certain conditions on the constraint store that must hold at certain points of program derivations. For concreteness we will use the `pred` assertions of the Ciao assertion language (Puebla et al. 1997; Hermenegildo et al. 1999; Puebla et al. 2000a,b; Hermenegildo et al. 2012). The main intent behind the construction of a specification for a predicate using `pred`-assertions is to define the set of all admissible preconditions for this predicate, and for each such precondition in turn specify the respective postcondition. I.e., `pred`-assertions allow stating sets of related *preconditions* and *conditional postconditions* for a given predicate.

These pre- and postconditions are formulas containing literals corresponding to predicates that are specially labeled as *properties*. The design of this language is such that properties and the other predicates composing the program are written in the same language. This approach is motivated by the direct correspondence between the declarative and operational semantics of constraint logic programs and it provides a direct link between the properties used in assertions and the corresponding run-time tests, which constitute (instrumented) calls to the predicates defining the properties. This also allows defining specifications that are more general than, e.g., the traditional notions of types.

More formally, the set of assertions for a given predicate represented by *Head* is composed of all statements of the form:<sup>1</sup>

$$\begin{aligned} & :- \text{pred } Head : Pre_1 \Rightarrow Post_1 . \\ & \dots \\ & :- \text{pred } Head : Pre_n \Rightarrow Post_n . \end{aligned}$$

where *Head* is the same normalized atom, that denotes the predicate that the assertions apply to, and the  $Pre_i$  and  $Post_i$  are conjunctions<sup>2</sup> of *prop* literals that refer to the variables of *Head*.

A set of assertions as above states that in any execution state  $\langle Head :: G \mid \theta \rangle$  at least one of the  $Pre_i$  conditions should hold, and that, given the  $(Pre_i, Post_i)$  pair(s) where  $Pre_i$  holds, then, if *Head* succeeds, the corresponding  $Post_i$  should hold upon success. More formally, given a predicate represented by a normalized atom *Head*, and the corresponding set of assertions is  $\mathcal{A}(Head) = \{A_1 \dots A_n\}$ , with  $A_i = \text{“} :- \text{pred } Head : Pre_i \Rightarrow Post_i \text{”}$  such assertions are normalized into a set of *assertion conditions* for that predicate, denoted as  $\mathcal{A}_C(Head) = \{C_0, C_1, \dots, C_n\}$  s.t.:

$$C_i = \begin{cases} c_i.\text{calls}(Head, \bigvee_{j=1}^n Pre_j) & i = 0 \\ c_i.\text{success}(Head, Pre_i, Post_i) & i = 1..n \end{cases}$$

where  $c_i$  is a unique assertion condition identifier. If there are no assertions associated with *Head* then the corresponding set of assertion conditions is empty. The set of assertion conditions for a program is the union of the assertion conditions for each of the predicates in the program.

The  $\text{calls}(Head, \dots)$  conditions encode the checks that ensure that the calls to the predicate represented by the *Head* literal are within those admissible by the set of assertions, and we thus call them the *calls assertion conditions*. The  $\text{success}(Head_i, Pre_i, Post_i)$  conditions encode the checks for compliance of the successes for particular sets of calls, and we thus call them the *success assertion conditions*.

### 2.3 Semantics with assertions

We now recall the operational semantics with assertions, which checks whether assertion conditions hold or not while comput-

<sup>1</sup> We follow the more compact formalization of (Stulova et al. 2014), using only `pred` assertions. See also (Puebla et al. 2000b) for the original presentation using `calls` and `success` assertions. We are also not dealing herein with `comp` assertions and properties.

<sup>2</sup> In the general case  $Pre$  and  $Post$  can be DNF formulas of *prop* literals but we limit them to conjunctions herein for simplicity of presentation.

ing the derivations from a query. In order to keep track of any violated assertion conditions, we use the identifiers of the assertion conditions. Given the atom  $L_a$  that is a renaming of some normalized atom  $L$  s.t.  $L_a = \sigma(L)$  and the corresponding set of assertion conditions  $\mathcal{A}_C(L)$ , the assertion conditions for  $L_a$  are obtained as follows: if  $\exists C \in \mathcal{A}_C(L), C = c.\text{calls}(L, Pre)$  (or  $C = c.\text{success}(L, Pre, Post)$ ), then  $C_a = \sigma(C) = c_a.\text{calls}(L_a, \sigma(Pre))$  (or  $C_a = \sigma(C) = c_a.\text{success}(L_a, \sigma(Pre), \sigma(Post))$ ). We also introduce an extended program state of the form  $\langle G \mid \theta \mid \mathcal{E} \rangle$ , where  $\mathcal{E}$  denotes the set of identifiers for falsified assertion condition instances and  $|\mathcal{E}| \leq 1$ . For the sake of readability, we write labels in *negated* form when they appear in the error set. A finished derivation from a query  $(L, \theta)$  now is *successful* if the last state is of the form  $\langle \square \mid \theta' \mid \emptyset \rangle$  ( $\emptyset$  denotes the empty set), and *failed* if the last state is of the form  $\langle L' \mid \theta' \mid \{\bar{c}\} \rangle$ . We also extend the set of literals with syntactic objects of the form  $\langle L, c \rangle$  where  $L$  is a literal and  $c$  is an identifier for an assertion condition instance, which we call *check literals*. Thus, a *literal* is now a constraint, an atom or a check literal. A literal  $L$  *succeeds trivially* for  $\theta$  in program  $P$ , denoted  $\theta \Rightarrow_P L$ , iff  $\exists \theta' \in \text{answers}((L, \theta))$  such that  $\theta \models \theta'$ . We can now recall the notion of *Reductions in Programs with Assertions* from (Stulova et al. 2014), which is our starting point: a state  $S = \langle L :: G \mid \theta \mid \emptyset \rangle$ , where  $L$  is a literal, can be *reduced* to a state  $S'$ , denoted  $S \rightsquigarrow_A S'$ , as follows:

1. If  $L$  is a constraint then the new state is  $S' = \langle G' \mid \theta' \mid \emptyset \rangle$  where  $G'$  and  $\theta'$  are obtained in a same manner as in  $\langle L :: G \mid \theta \rangle \rightsquigarrow \langle G' \mid \theta' \rangle$
2. If  $L$  is an atom and  $\exists(L :- B) \in \text{defn}(L)$ , then the new state  $S'$  is obtained as

$$S' = \begin{cases} \langle L \mid \theta \mid \{\bar{c}\} \rangle & \text{if } \exists c.\text{calls}(L, Pre) \in \mathcal{A}_C(L) \\ & \wedge \theta \not\Rightarrow_P Pre \\ \langle B :: G' \mid \theta \mid \emptyset \rangle & \text{otherwise} \end{cases}$$

and  $G' = \text{acheck}(L, c_1) :: \dots :: \text{acheck}(L, c_n) :: G$  such that  $c_i.\text{success}(L, Pre_i, Post_i) \in \mathcal{A}_C(L) \wedge \theta \Rightarrow_P Pre_i$ .

3. If  $L$  is a check literal  $\text{acheck}(L', c)$ , then  $S'$  is obtained as

$$S' = \begin{cases} \langle L' \mid \theta \mid \{\bar{c}\} \rangle & \text{if } c.\text{success}(L', \_, Post) \in \mathcal{A}_C(L') \\ & \wedge \theta \not\Rightarrow_P Post \\ \langle G \mid \theta \mid \emptyset \rangle & \text{otherwise} \end{cases}$$

## 3. Assertion Checking Modes

When a program is being instrumented with run-time checks the choice of instrumentation strategy is determined by several factors and considerations. Most of these factors can typically be generalized to a compromise between thoroughness of the code annotation (complexity of the properties, annotation depth) and the resulting performance penalties (execution time slowdown, code bloat, increases in memory use).

We propose a view on this compromise that differentiates between various levels of behavioral safety guarantees embodied in several *assertion checking modes*. In the following we will explain these modes for concreteness in a client-library interaction scenario, and from the perspective of the client. However, the proposed view is not limited to this particular case and can be applied to other kinds of interaction between program components.

**Unsafe Checking Mode** This checking mode corresponds to a scenario where no execution time slowdown is tolerated at run time, even at the cost of no program behavior safety guarantees. In this case no run-time checks are generated from the assertions of the library. Formally, this corresponds to using the standard semantics of Section 2.1, and thus ignoring the assertions in the program. This of course eliminates any overhead but at the cost of correctness.

```

1 :- module(_, [p]). % exported predicate p
2
3 :- check pred p : Pre => Post.
4
5 % c0.calls(p, Pre)          ^ status(c0, check)
6 % c1.success(p, Pre, Post) ^ status(c1, check)
7
8 p :- body.
9
10 q :- p.

```

(a) Initial program fragment.

```

1 :- module(_, [p]).
2
3 % c0.calls(p, Pre)          ^ status(c0, check)
4 % c1.success(p, Pre, Post) ^ status(c1, check)
5
6 p :- p_inner. % (the link clause)
7
8 p_inner :- body.
9
10 q :- p_inner.

```

(b) The same program fragment after the transformation.

**Figure 1.** Client-safe program transformation.

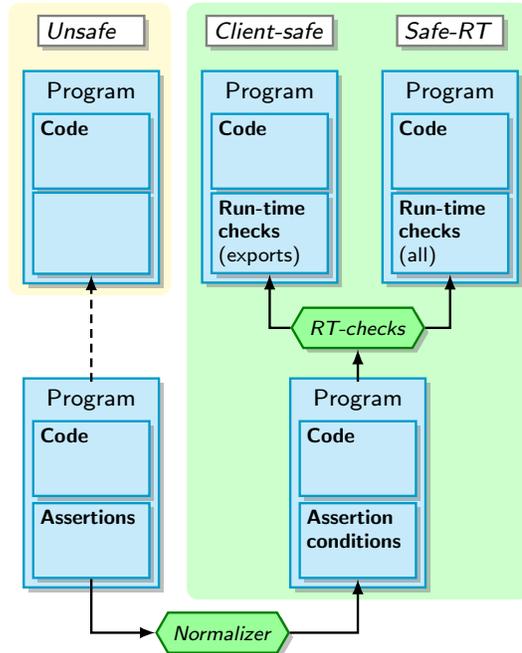
However, we still consider it because it represents a baseline and is in fact used sometimes in practice for production code, in order to avoid overhead, if, e.g., it is perceived that sufficient testing was carried out prior to delivery.

**Client-safe Checking Mode** In this mode the library provides the client with behavior guarantees on its interface, but does not check the assertions for the internal procedures. Run-time checks are thus generated only for the assertion conditions for the exported predicates of the library. More formally, assuming that the set of (atoms of) exported predicates is given by  $Exp$ :

1. If  $L$  is a constraint or  $L$  is an atom such that  $L \notin Exp$ , then the new state  $S' = \langle G' \mid \theta' \mid \emptyset \rangle$  where  $G'$  and  $\theta'$  are obtained in a same manner as in  $\langle L :: G \mid \theta \rangle \rightsquigarrow \langle G' \mid \theta' \rangle$ .
2. If  $L$  is an atom such that  $L \in Exp$ , and  $\exists(L :- B) \in \text{defn}(L)$ , then the new state  $S'$  is obtained as in step 2 of the semantics with run-time checks of Section 2.3.
3. If  $L$  is a check literal  $\text{acheck}(L', c)$ , then  $S'$  is obtained as in step 3 of the semantics with run-time checks of Section 2.3.

The modified semantics above ensures that checks are performed only for the predicates in the library interface. However, all calls within the library to the exported predicates, including recursive calls, would also be checked, unnecessarily. In order to avoid this, and to ensure that the checks are performed only on the external calls, we assume that the program transformation given in Fig. 1 is applied to all exported predicates. The transformation introduces intermediate *link* predicates for exported predicates so that the module interface is preserved but all the internal calls are replaced by calls to the wrapper predicates, for which no checks are performed. This combination of program transformation and run-time checking policy allows obtaining safety guarantees at the library boundaries with minimal run-time checking execution time overhead.

**Safe-RT Execution Mode** In this mode the library provides behavior guarantees both on its interface and its internals. Run-time checks are generated for all assertions of the library. This corresponds to using the semantics with assertions of Section 2.3. The performance penalty here is the largest.



**Figure 2.** Source transformation differences per checking mode.

**Transformations** The checking modes described above require different source transformations to be performed on a program during compile time (see Fig. 2). Before any such transformations take place, the assertions are normalized and expanded into assertion conditions. This allows assuring that no syntactic errors are present in the assertion conditions and that no undefined properties (i.e., properties that are not defined in the program or imported from libraries) appear in such conditions.

In the *Unsafe* mode nothing is done and the assertion conditions are simply ignored during compilation. In the *Safe-RT* mode the source transformation is quite straightforward: all the assertion conditions for all assertions in the program are turned into run-time checks directly. In the *Client-safe* mode, as mentioned before, the program transformation of Figure 1 is first performed for all the exported predicates, and then run-time checks are generated only for the assertion conditions of those exported predicates.

## 4. Optimizing Run-Time Checks via Static Analysis

We now return to the issue of optimizing run-time checks via (abstract interpretation-based) static program analysis, in order to reduce the number of run-time tests and thus the overhead from run-time testing, following the Ciao model. To this end, we recall the basic abstract interpretation-based analysis approach used and the memo table representation of the analysis results and describe how run-time tests are optimized using the information in the analysis memo table. Based on this in the following section we will present our approach for taking advantage of the run-time checking semantics to improve the precision of the analysis.

Herein we will refer to this combination of static and dynamic checking as the **Safe-CT-RT Checking Mode**, i.e., as a variation on the *Safe-RT* run-time checking mode, where static verification is performed in order to eliminate as many of the properties in the program assertions to be checked at run time as possible. Run-time checks are still generated for all program assertions but in contrast to the *Safe-RT* case the assertions are simplified before the

checks are generated from them. In this mode the run-time checks for the *calls* assertion conditions of the exported predicates are left untouched though, to ensure calls safety in an open context.<sup>3</sup>

#### 4.1 Abstract Interpretation-based Analysis

For analysis we use the technique of abstract interpretation (Cousot and Cousot 1977), which simulates the execution of a program on an *abstract domain* ( $D_\alpha$ ) which is simpler than the actual, *concrete domain*<sup>4</sup> ( $D$ ). Abstract values and sets of concrete values are related via a pair of monotonic mappings  $\langle \alpha, \gamma \rangle$ : *abstraction*  $\alpha : 2^D \rightarrow D_\alpha$ , and *concretization*  $\gamma : D_\alpha \rightarrow 2^D$ . The operations of *least upper bound* ( $\sqcup$ ) and *greatest lower bound* ( $\sqcap$ ) over abstract literals  $\lambda$  mimic those of  $2^D$  in a precise sense:

$$\begin{aligned} \forall \lambda, \lambda' \in D_\alpha : \lambda \sqsubseteq \lambda' &\Leftrightarrow \gamma(\lambda) \subseteq \gamma(\lambda') \\ \forall \lambda_1, \lambda_2, \lambda' \in D_\alpha : \lambda_1 \sqcup \lambda_2 = \lambda' &\Leftrightarrow \gamma(\lambda_1) \cup \gamma(\lambda_2) = \gamma(\lambda') \\ \forall \lambda_1, \lambda_2, \lambda' \in D_\alpha : \lambda_1 \sqcap \lambda_2 = \lambda' &\Leftrightarrow \gamma(\lambda_1) \cap \gamma(\lambda_2) = \gamma(\lambda') \end{aligned}$$

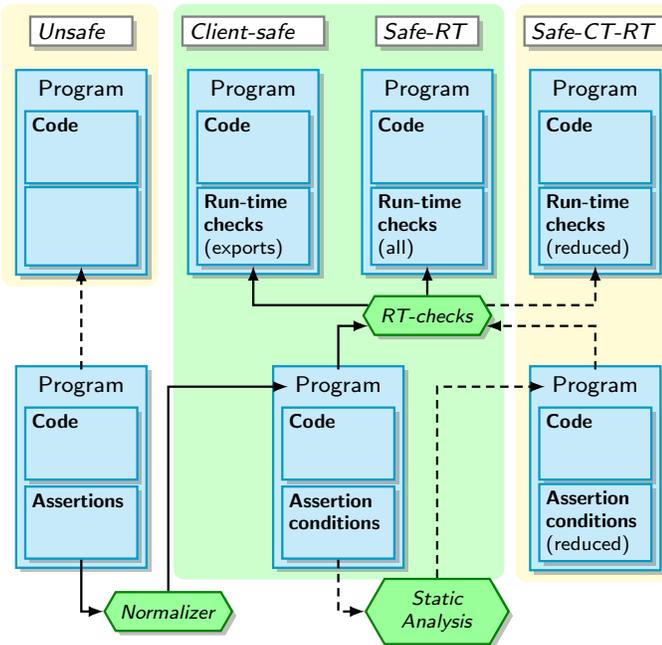
As usual in abstract interpretation,  $\perp$  denotes the abstract constraint such that  $\gamma(\perp) = \emptyset$ , whereas  $\top$  denotes the most general abstract constraint, i.e.,  $\gamma(\top) = D$ .

The concrete framework that we will use in the static analysis component is the Ciao PLAI abstract interpretation system (Muthukumar and Hermenegildo 1989, 1990, 1992). Below we adapt some definitions and notation from (Puebla et al. 2000b) to illustrate the analysis process implemented by PLAI.

The goal dependent abstract interpretation performed by PLAI takes as input a program  $P$ , an abstract domain  $D_\alpha$ , and a description  $\mathcal{Q}_\alpha$  of the possible initial queries to the program given as a set of abstract queries. An *abstract query* is a pair  $(L, \lambda)$ , where  $L$  is an atom (for one of the exported predicates) and  $\lambda \in D_\alpha$  an

<sup>3</sup>Although modular analysis can also eliminate the module interface checks, this is to be consistent with our simple library scenario.

<sup>4</sup>In what follows we assume the concrete domains to have a powerset structure, but the framework is not limited to such domains and can be applied to domains of arbitrary structure.



**Figure 3.** Source transformation differences per checking mode, including compile-time analysis.

**Table 1.** Assertion status.

Status	Source	Description
check	user	The assertion expresses part of the intended semantics. It may or may not hold in the current version of the program. It is the default status that is assumed for assertions written without and explicit status.
checked	static checking	The assertion was a check assertion which has been proved to actually hold in the current version of the program for any valid initial call.
false	static checking	Similarly, a check assertion is rewritten with the status false when it is proved not to hold for some valid initial query.
true	static analyses or user	Such an assertion expresses (a part of) the actual semantics of the program, normally automatically inferred by analysis. Can also be provided by the programmer.

abstract constraint which describes the initial stores for  $L$ . A set of abstract queries  $\mathcal{Q}_\alpha$  represents a set of queries, denoted  $\gamma(\mathcal{Q}_\alpha)$ , which is defined as  $\gamma(\mathcal{Q}_\alpha) = \{(L, \theta) \mid (L, \lambda) \in \mathcal{Q}_\alpha \wedge \theta \in \gamma(\lambda)\}$ . Such an abstract interpretation computes a set of triples  $Analysis(P, \mathcal{Q}_\alpha, D_\alpha) = \{\langle L_p, \lambda^c, \lambda^s \rangle \mid p \text{ is a predicate of } P\}$ , where  $\lambda^c$  and  $\lambda^s$  are abstract constraints that describe calls and success patterns for  $p$ .

The analysis (as the assertion language, to be introduced later) is designed to discern among the various usages of a predicate. Thus, multiple usages of a procedure can result in multiple descriptions in the analysis output, i.e., for a given predicate  $P$  multiple  $\langle P, \lambda^c, \lambda^s \rangle$  triples may be inferred. More precisely, the analysis is said to be *multivariant on calls* if more than one triple  $\langle P, \lambda_1^c, \lambda_1^s \rangle, \dots, \langle P, \lambda_n^c, \lambda_n^s \rangle$   $n \geq 0$  with  $\lambda_i^c \neq \lambda_j^c$  for some  $i, j$  may be computed for the same predicate. However, for simplicity of presentation, we assume that the analysis computes exactly one tuple  $\langle L_p, \lambda^c, \lambda^s \rangle$  for each (reachable) predicate  $p$ .

#### 4.2 Optimizing Assertions with Analysis Results

The steps of the verification process are represented by associating a notion of “status” to each assertion:

$$\begin{aligned} &:- [Status] \text{ pred Head} : Pre_1 \Rightarrow Post_1. \\ &\dots \\ &:- [Status] \text{ pred Head} : Pre_n \Rightarrow Post_n. \end{aligned}$$

This optional *Status* flag indicates whether the assertion refers to intended or actual properties, and possibly some additional information, as shown in the top part of Table 1 (see also Figure 6).

The reasoning about the statuses of assertion conditions is performed in the following terms. Given a literal  $L$  and a program  $P$ , the *trivial success set* of  $L$  in  $P$  is  $TS(L, P) = \{\exists_L \theta \mid \theta \Rightarrow_P L\}$ . We also recall here the auxiliary partial functions *prestep* and *step* from (Stulova et al. 2014) which are instrumental in reasoning about program state reductions:

$$\begin{aligned} \text{prestep}(L_a, D) &= (\theta, \sigma) \equiv D_{[-1]} = \langle L :: G \mid \theta \rangle \wedge \exists \sigma L = \sigma(L_a) \\ \text{step}(L_a, D) &= (\theta, \sigma, \theta') \equiv D_{[-1]} = \langle G \mid \theta' \rangle \wedge \exists \sigma L = \sigma(L_a) \\ &\quad \wedge \exists i D_{[i]} = \langle L :: G \mid \theta \rangle \end{aligned}$$

Given a derivation whose current state is a call to  $L_a$  (normalized atom), the *prestep* function returns the substitution  $\sigma$  for  $L_a$ , and the constraint store  $\theta$  at the predicate *call* (i.e., just before the literal is reduced). Given a derivation whose current state corresponds

exactly to the return from a call to  $L_a$ , the step function returns the substitution  $\sigma$  for  $L_a$ , the constraint store  $\theta$  at the call to  $L_a$ , and the constraint store  $\theta'$  at  $L_a$ 's *success* (i.e., just after all literals introduced from the body of  $L_a$  have been fully reduced).

An abstract constraint  $\lambda_{TS(L,P)}^-$  is an *abstract trivial success subset* of  $L$  in  $P$  iff  $\gamma(\lambda_{TS(L,P)}^-) \subseteq TS(L,P)$ . An abstract constraint  $\lambda_{TS(L,P)}^+$  is an *abstract trivial success superset* of  $L$  in  $P$  iff  $\gamma(\lambda_{TS(L,P)}^+) \supseteq TS(L,P)$ . Given the program  $P$ , the concrete and abstract sets of queries  $\mathcal{Q}$  and  $\mathcal{Q}_\alpha$ <sup>5</sup> respectively, where  $\gamma(\mathcal{Q}_\alpha) \supseteq \mathcal{Q}$ , and  $\langle L, \lambda^c, \lambda^s \rangle \in \text{Analysis}(P, \mathcal{Q}_\alpha, D_\alpha)$ , the status of an assertion condition  $C$ , associated with it by the mapping  $\text{status}(c, \text{Status})$  where  $c$  is the corresponding identifier, is determined as follows:

- $C = c.\text{calls}(L, \text{Precond}) \wedge \text{status}(c, \text{checked})$   
if  $\lambda^c \sqsubseteq \lambda_{TS(\text{Precond}, P)}^-$ .
- $C = c.\text{success}(L, \text{Pre}, \text{Post}) \wedge \text{status}(c, \text{checked})$   
if (1)  $\lambda^c \sqcap \lambda_{TS(\text{Pre}, P)}^+ = \perp$  or (2)  $\lambda^s \sqsubseteq \lambda_{TS(\text{Post}, P)}^-$ ;
- $C = c.\text{calls}(L, \text{Precond}) \wedge \text{status}(c, \text{false})$   
if  $\exists D \in \text{derivs}(\mathcal{Q})$  s.t.  $\text{prestep}(L, D) = (\theta, \sigma) \wedge \bar{\lambda}_L \theta \neq \emptyset$  and  $\lambda^c \sqcap \lambda_{TS(\text{Precond}, P)}^+ = \perp$ .
- $C = c.\text{success}(L, \text{Pre}, \text{Post}) \wedge \text{status}(c, \text{false})$   
if  $\lambda^c \sqcap \lambda_{TS(\text{Pre}, P)}^- \neq \perp$  and  $\lambda^s \sqcap \lambda_{TS(\text{Post}, P)}^+ = \perp$  and  $\exists \theta \in \gamma(\lambda^c \sqcap \lambda_{TS(\text{Pre}, P)}^-) : \exists D \in \text{derivs}(\mathcal{Q})$  s.t.  $\text{step}(L, D) = (\theta, \sigma, \theta') \wedge \bar{\lambda}_L \theta' \neq \emptyset$ .

The compile-time checking process can be seen as a revision of the assertion statuses where for each predicate literal  $L$  its *annotation* composed from the respective assertion conditions  $\mathcal{A}_C^{usr}(L) = \{C \mid c.C \in \mathcal{A}_C(L) \wedge \text{status}(c, S) \wedge S \in \{\text{check}, \text{true}\}\}$  given the analysis output of the form  $\mathcal{A}_C^{ana}(L) = \{C \mid \forall c.C \text{ status}(c, \text{true})\}$  is rewritten into  $\{C \mid c.C \in \mathcal{A}_C^{usr}(L) \wedge \text{status}(c, S) \wedge S \in \{\text{check}, \text{checked}, \text{false}\}\}$ .

## 5. Taking Advantage of the Run-Time Checking Semantics during Analysis

The standard analysis introduced in Section 4.1 safely approximates the traditional semantics (i.e., the semantics without assertions or run-time checks).<sup>6</sup> However, if we know that run-time checks will be performed for sure for a certain set of (check) assertions (as, e.g., for all assertions in the Safe-RT execution mode, or the ones corresponding to interface predicates in the client-safe mode), it is possible to use this information during analysis to improve precision, i.e., to assume that the calls assertion conditions hold after the predicate has entered the predicate definition (since, according to the semantics of Section 2.3 either the checks for these calls assertion conditions have already succeeded or the program has exited with error), and to assume the relevant success assertion conditions after the predicate has exited (since, again, at this point

<sup>5</sup> In the implementation of PLAI,  $\mathcal{Q}_\alpha$  is obtained from the calls conditions of the assertions of exported predicates (or, if no such assertions are present, a “topmost” abstract state is assumed), or from specific “entry” assertions.

<sup>6</sup> Actually, assertions with `trust` status (which are typically used to provide information to improve analysis precision), and assertions with `true` status are in fact read and applied by the traditional analysis during its fixpoint calculation. However, in this discussion we refer to incorporating into the analysis the information present in `check` assertions, i.e., from the assertions being checked at compile time or run time. These assertions are not normally taken into account by the analysis since they may or may not hold and, in general, run-time tests may or may not be included in the compiled program.

```

1 :- module(_, [p/2]).           % p/2 is exported
2 :- use_module(foo, [e/2]).    % e/2 is imported
3
4 p(X,Y) :- q(X,Y).
5
6 :- pred q(X,Y) : int(X) => int(Y).
7 q(X,Y) :- r(X,Y).
8
9 :- pred r(X,Y) : int(X) => int(Y).
10 r(X,Y) :- e(X,Y).

```

Figure 4. Example for analysis improvement.

either these success assertion conditions have already succeeded or the program has exited with error).

As an example, consider the program of Figure 4.<sup>7</sup> `p/2` is an exported predicate, `q/2` and `r/2` are local predicates, and `e/2` is imported. We allow both `p/2` and `e/2` to be called without any restriction, and we do not specify any constraints either regarding their successes. However, we want to enforce (through the two assertions) that `q/2` and `r/2` always be called with their first argument  $X$  bound to an integer, and that their second argument  $Y$  be bound to an integer upon success. Since any type of call is allowed to `p/2`, and without information on the presence of run-time checks, analysis cannot infer anything about the calls conditions for `q/2` and `r/2`, or for the success conditions of these two predicates, and will report warnings for unchecked conditions for all of them (and the two assertions will remain in `check` status).

However, note that, assuming we are generating run-time checks for all assertion conditions, the call to `r/2` in the body of `q/2` can only be reached if the calls condition for `q/2` holds, i.e., if  $X$  is bound to an integer (since otherwise execution would have been aborted). Thus, this information can be incorporated into the analysis and propagated to the call to `r/2`, and it can be determined that the calls condition for `r/2` (i.e., that its first argument be also bound to an integer) always holds. Thus, this calls condition for `r/2` gets status `checked` and no run-time test needs to be generated for it.

Similarly, the run-time test for the success condition for `r/2` ensures that if the call to `r/2` in the body of `q/2` returns, then its second argument is guaranteed to be bound to an integer. Thus, the success condition for `q/2` will also get status `checked` and no run-time test needs to be generated for it either.

**Transformation** A straightforward method to incorporate the information from successful checks into the analysis, so that it takes the semantics with run-time checking into account, would be to analyze the transformed program (i.e., the program including the code that performs the run-time tests) instead of the original one. This is the approach implied by the original transformational definitions of the assertion language. On the other hand, programs transformed for run-time testing contain numerous optimizations and instrumentation that make their analysis less efficient and can potentially affect precision. An alternative would be to use a very simple (even if inefficient) run-time checking transformation just for analysis. Inspired by this idea, we propose herein a different, even more direct approach, based on introducing additional assertions and link predicates in the program that together capture the run-time checking semantics and provide the additional information source for the analysis, in order to increase precision. This is performed as a program transformation  $T$  that precedes the analysis and is applied to every annotated predicate in a program:

<sup>7</sup> In the examples we use just simple types as properties for conciseness, but even in this case please note that the use of types is *moded*, i.e., the assertions here express *states of instantiation*.

```

1 :- check calls   q(X,Y) : int(X).
2 :- true  success q(X,Y) : int(X) => int(Y).
3 q(X,Y) :- q_inner(X,Y).
4
5 :- true  calls   q_inner(X,Y) : int(X).
6 :- check success q_inner(X,Y) : int(X) => int(Y).
7 q_inner(X,Y) :- r(X,Y).

```

**Figure 5.** CTRT program transformation example (output).

$$T(L) = \{\{L:-L_{inner}\} \cup \text{defn}(L_{inner}), \mathcal{A}_C^{link} \cup \mathcal{A}_C^{inner}\}$$

where  $L = p(\vec{X})$ , and the literal  $L_{inner} = p_{inner}(\vec{X})$  is obtained with a new predicate symbol  $p_{inner}$ , and:

$$\begin{aligned} \text{defn}(L_{inner}) &= \{L_{inner}:-B \mid L:-B \in \text{defn}(L)\} \\ \mathcal{C} &= \{c.C \in \mathcal{A}_C(L) \mid \text{status}(c, \text{check})\} \\ \mathcal{A}_C^{link} &= \{c^l.C \mid c.C \in \mathcal{C}\} \text{ and } \forall c^l.C \in \mathcal{A}_C^{link} \text{ we extend} \\ &\text{the status relation s.t. } \text{status}(c^l, S^l), \text{ where:} \\ S^l &= \begin{cases} \text{check} & \text{if } C = \text{calls}(-, -) \\ \text{true} & \text{if } C = \text{success}(-, -, -) \end{cases} \\ \mathcal{A}_C^{inner} &= \{c^i.C \mid c.C \in \mathcal{C}\} \text{ and } \forall c^i.C \in \mathcal{A}_C^{inner} \text{ we extend} \\ &\text{the status relation s.t. } \text{status}(c^i, S^i), \text{ where:} \\ S^i &= \begin{cases} \text{true} & \text{if } C = \text{calls}(-, -) \\ \text{check} & \text{if } C = \text{success}(-, -, -) \end{cases} \end{aligned}$$

The objective of the transformation is to improve the precision and reduce the cost of the analysis, while preserving program behavior when the `check` assertion conditions are expanded into run-time checks. The transformation modifies all predicates with `check` assertions. For each such predicate  $p$ , the original predicate symbol is renamed into  $p_{inner}$  and a single-clause wrapper predicate for  $p$  (which we will refer to as a *link* clause), is introduced which calls the  $p_{inner}$  predicate.

The set of assertion conditions for the initial predicate  $p$  is duplicated for the  $p_{inner}$  counterpart, including their original statuses. However, the statuses of the `success` assertion conditions for  $p$  in the link clause and the `calls` assertion conditions of  $p_{inner}$  are set to `true`. As a result, the `calls` assertion conditions for  $p$  (i.e.,  $c^l.\text{calls}(L, -)$  with  $\text{status}(c^l, \text{true})$ ) will still be checked in the version with run-time checks, but they will be assumed in  $p_{inner}$  (i.e.,  $c^i.\text{calls}(L_{inner}, -)$  with  $\text{status}(c^i, \text{true})$ ).

For the success part the assertion conditions will still be checked for the inner predicate (i.e.,  $c^i.\text{success}(L_{inner}, -, -)$  with  $\text{status}(c^i, \text{check})$ ) and the information will be assumed upon exiting  $p$  (i.e.,  $c^i.\text{success}(L, -, -)$  with  $\text{status}(c^i, \text{true})$ ). The transformation guarantees that the same run-time tests will be performed, that no duplication of checks will occur (since there are no intermediate states between the calls to  $p$  and  $p_{inner}$  and exits from  $p_{inner}$  to  $p$ ), and that the analysis will gather the right information.

An example of the CTRT transformation for the `q/2` predicate from the program in Fig. 4 is shown in Fig. 5. The `true` assertions here correspond to the additional information that can be safely used in the analysis. Since all predicates with assertions undergo this transformation, a number of inner calls coming from the link clauses is added to the program. Yet such calls are relatively inexpensive and the resulting runtime overhead is negligible. Even more, should the analysis verify the `calls` assertion condition of the link clause or the `success` assertion condition of the inner clause, the link clause then becomes unnecessary and can be completely removed.

**Lemma 1** (Correctness of the CTRT Transformation) *Let  $P$  be a program and  $Q = (L, \theta)$  a query to  $P$ . Then  $\forall D \in \text{derivs}(Q)$  the final state  $D_{[-1]}$  is the same in the versions of  $P$  with and without the CTRT transformation.*

*Proof.* First, let us prove the correctness of the transformation for the `calls` assertion conditions.

Let  $\mathcal{A}_C(L) = \{C\}$  where  $C = c.\text{calls}(L, Pre)$  s.t.  $\text{status}(c, \text{check})$  and  $\exists(L:-B) \in \text{defn}(L)$ . The possible reduction sequences from the  $S_0 = \langle L :: G \mid \theta \mid \emptyset \rangle$  state are:

$$\begin{aligned} S_0 &\rightsquigarrow_{\mathcal{A}} \langle B :: G \mid \theta \mid \emptyset \rangle = S_{succ} \quad \text{if } \theta \Rightarrow_P Pre \\ S_0 &\rightsquigarrow_{\mathcal{A}} \langle L \mid \theta \mid \{\bar{c}\} \rangle = S_{err} \quad \text{if } \theta \not\Rightarrow_P Pre \end{aligned}$$

Now let us add the *link* clause for  $L$  and rename its other clauses s.t.  $\text{defn}(L) = \{L:-L_{inner}\}$  and  $\exists L_{inner}:-B \in \text{defn}(L_{inner})$ , and let's add an assertion condition for  $L_{inner}$ :  $C^{inner} = c^i.\text{calls}(L_{inner}, Pre)$  with  $\text{status}(c^i, \text{check})$ . The possible reduction sequences from the  $S_0$  state now are:

$$\begin{aligned} S_0 &\rightsquigarrow_{\mathcal{A}} \langle L_{inner} :: G \mid \theta \mid \emptyset \rangle \rightsquigarrow_{\mathcal{A}} S_{succ} \quad \text{if } \theta \Rightarrow_P Pre \\ S_0 &\rightsquigarrow_{\mathcal{A}} S_{err} \quad \text{if } \theta \not\Rightarrow_P Pre \end{aligned}$$

The  $S_0 \rightsquigarrow_{\mathcal{A}} \langle L_{inner} :: G \mid \theta \mid \emptyset \rangle \rightsquigarrow_{\mathcal{A}} \langle L_{inner} \mid \theta \mid \{\bar{c}^i\} \rangle$  reduction sequence is impossible since it would require  $\theta \Rightarrow_P Pre$  to hold in the first reduction step and  $\theta \not\Rightarrow_P Pre$  to hold in the second reduction step.

This way in both assertion checking modes  $D_{[-1]} \in \{S_{succ}, S_{err}\}$  and run-time checks for the `calls` assertion condition  $C^{inner}$  (namely, checks for  $\theta \Rightarrow_P Pre$  after the checks for  $\theta \Rightarrow_P Pre$ ) could be safely removed by setting  $\text{status}(c^i, \text{true})$ .

Next, let's consider the case of `success` assertion conditions.

Let  $\mathcal{A}_C(L) = \{C\}$  where  $C = c.\text{success}(L, Pre, Post)$  s.t.  $\text{status}(c, \text{check})$  and  $\exists(L:-B) \in \text{defn}(L)$ . The possible reduction sequences from the  $S_0 = \langle L :: G \mid \theta \mid \emptyset \rangle$  state are:

$$\begin{aligned} S_0 &\rightsquigarrow_{\mathcal{A}} \langle B :: \text{acheck}(L, c) :: G \mid \theta \mid \emptyset \rangle \rightsquigarrow_{\mathcal{A}}^* \langle G \mid \theta \mid \emptyset \rangle = S_{succ} \\ &\quad \text{if } \theta \Rightarrow_P Post \\ S_0 &\rightsquigarrow_{\mathcal{A}} \langle B :: \text{acheck}(L, c) :: G \mid \theta \mid \emptyset \rangle \rightsquigarrow_{\mathcal{A}}^* \\ &\quad \langle \text{acheck}(L, c) \mid \theta \mid \bar{c} \rangle = S_{err} \\ &\quad \text{if } \theta \not\Rightarrow_P Post \end{aligned}$$

Now let us add the *link* clause for  $L$  and rename its other clauses s.t.  $\text{defn}(L) = \{L:-L_{inner}\}$  and  $\exists L_{inner}:-B \in \text{defn}(L_{inner})$ , and let's add an assertion condition for  $L_{inner}$ :  $C^{inner} = c^i.\text{success}(L_{inner}, Pre, Post)$  with  $\text{status}(c^i, \text{check})$ . We also now consider  $C$  as  $C^{link}$  with its identifier  $c^l$ . The possible reduction sequences from the  $S_0$  state now are:

$$\begin{aligned} S_0 &\rightsquigarrow_{\mathcal{A}} \langle B :: \text{acheck}(L, c^i) :: \text{acheck}(L, c^l) :: G \mid \theta \mid \emptyset \rangle \rightsquigarrow_{\mathcal{A}}^* \\ &\quad \langle G \mid \theta \mid \emptyset \rangle = S_{succ} \\ &\quad \text{if } \theta \Rightarrow_P Post \\ S_0 &\rightsquigarrow_{\mathcal{A}} \langle B :: \text{acheck}(L, c^i) :: \text{acheck}(L, c^l) :: G \mid \theta \mid \emptyset \rangle \rightsquigarrow_{\mathcal{A}}^* \\ &\quad \langle \text{acheck}(L, c^l) \mid \theta \mid \bar{c}^l \rangle = S_{err} \\ &\quad \text{if } \theta \not\Rightarrow_P Post \end{aligned}$$

Although the assertion condition identifiers for the two  $S_{err}$  are different, the checks performed in these states are equal ( $\theta \not\Rightarrow_P Post$ ).

This way the run-time checks for the  $c^l$  assertion condition are duplicating the checks for  $c^i$  and could be safely removed by setting  $\text{status}(c^l, \text{true})$ .  $\square$

## 6. Experiments

As stated throughout the paper, our objective is to explore the effectiveness of abstract interpretation in detecting parts of program specifications that can be statically simplified to true or false, and to quantify the impact of this application of analysis towards reducing the cost of the run-time checks. In particular, we have studied these issues for the different assertion checking modes that we have defined.

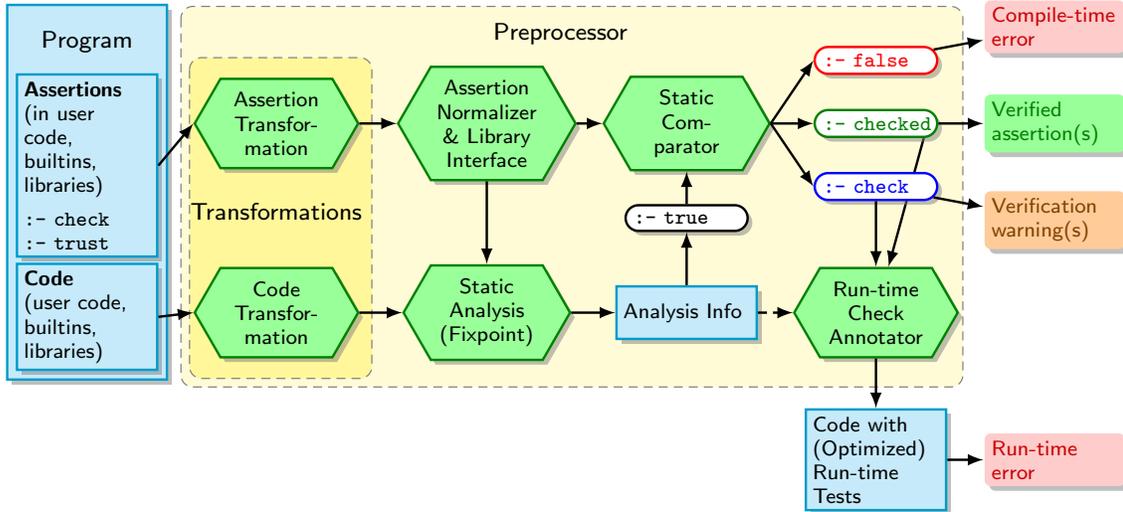


Figure 6. Adding the transformations to the Ciao Preprocessor.

## 6.1 Experimental Setup

We have built an experimental harness by extending the Ciao preprocessor, CiaoPP, which implements our baseline assertion verification framework. The architecture of this framework is shown in Figure 6. In that figure, hexagons represent system tools and components and arrows indicate the communication paths among them. Most of this communication is performed also in terms of assertions.

The input to the verification process is the user program, optionally including a set of assertions; this set always includes any assertions present for predicates exported by any libraries used. The `check` and `trust/true` assertions are normalized and the program is expanded to kernel form (simple horn clauses), and they are all given as input to the *static analysis*.

We have introduced new front-end passes implementing the new transformations (marked in Figure 6) which thus support the defined scenarios, as well as some other minor adaptations and extensions to the interface to select these different scenarios.

The results of analysis over the different abstract domains selected are provided in the form of `true` assertions. Then, for every predicate  $p$  in the program the framework performs compile-time checking of assertions by *comparing* the check assertions in the program (their assertion conditions) with the analysis results.

As a possible result of the comparison, assertions may be proved to hold, in which case they get `checked` status. As another possible result, assertions can be proved not to hold, in which case they get `false` status and a compile-time error is issued. Finally, if it is not possible to prove nor to disprove (part of) an assertion, then such assertion (or the relevant subset) is left as a `check` assertion, and the run-time check annotator introduces run-time checking code in the program for the assertion conditions as required by the scenario. In particular, the program transformations used in our experiments are those of (Stulova et al. 2015), with no caching.

## 6.2 Properties and Analysis Domains

In our experiments we concentrate on two classes of properties. The first one is the *state of variable instantiation*, i.e., which variables are bound to ground terms, or unbound, and, if they are unbound, the sharing (aliasing) patterns in order to be able to transfer accurately grounding information (“strong update”). These properties are approximated safely and quite accurately using the *sharing and freeness* domain (Muthukumar and Hermenegildo 1991). The sec-

ond class of properties we will be using refers to the shapes of the data structures constructed by the program in memory. To this end we use the *eterms* (Vaucheret and Bueno 2002) abstract domain which infers safely these shapes as regular trees.<sup>8</sup>

## 6.3 Benchmarks

To study the differences in the run-time overhead levels observed in different assertion checking modes we have selected a set of benchmarks, listed in Table 2.<sup>9</sup> These benchmarks represent simple yet diverse programs that represent frequently-occurring programming patterns such as performing symbolic or arithmetic computations, problem solving in fixed domains, processing stream data, etc. In general, they include recursion, search, irregular/dynamic data structures, etc. The relative internal complexity despite their generally small size make them good candidates to answer our main questions, allowing us to concentrate on the properties of interest in each case. All the benchmarks have been carefully annotated with reasonable program assertions that describe the expected behaviour. E.g., this is a selection of the `fft` code:

```

1 :- regtype complex/1. % A complex number
2 complex((A,B)) :- num(A), num(B).
3

```

<sup>8</sup>Note that these notions are more general and powerful than the traditional notions of types of modes. Also, comparing to the *traditional* notion of type inference in statically-typed languages, not only the types are generalized to any property supported by an abstract domain, but also the overall approach is quite different: there all the type definitions must be present in the program, and the inference problem just amounts to assigning one of these types to each program element in a single pass. If this assignment is not possible the program is rejected. In the Ciao approach, no type definitions are required. The purpose of analysis is precisely to infer, in a closed form (e.g., regular types) the shapes of the data structures that are built in memory for the whole program, which is done via a fixpoint calculation. Also, note that the regular types inferred and checked allow sub-typing. The situation is similar for the sharing+freeness domain versus, for example, traditional modes. This is also a strong difference with other approaches within logic programming, such as Mercury or Gödel, which also require the type definitions to be provided and that the program be typeable. See (Hermenegildo et al. 2011) for a further discussion to the very interesting topic of how to best straddle the dynamic vs. static language boundaries.

<sup>9</sup>Source available at <https://cliplab.org/papers/optchk-pdp2016/>

```

4 :- pred complex_mul(A, B, C) % Multiplication
5   : complex * complex * term
6   => complex * complex * complex.
7 complex_mul((Ra,Ia), (Rb,Ib), (Rs,Is)) :-
8   Rs is Ra*Rb-Ia*Ib,
9   Is is Ra*Ib+Rb*Ia.

```

Table 3 presents some quantitative characteristics of the benchmarks, such as lines of code (LOC) and size metrics, and also the total number of program assertions.

While the effectiveness of our assertion-based approach is not directly the objective of this paper, it is worth noticing that during our experiments the analysis on one of the more complex programs, *boyer*, has allowed us to spot bugs in the original translation from LISP that had been around for 30 years.

#### 6.4 Experimental Results

Table 4 shows the compilation time for the benchmarks under the different assertion checking modes.<sup>10</sup> The compilation time for the benchmarks under the *Safe-CT-RT* mode includes the total static

<sup>10</sup> Times for compilation and analysis assume that the compiler and analyzer are already loaded in memory and ready to execute. Thus, we removed the compiler and CiaoPP start-up time. In the current implementation, the engine needs around 1.4 seconds to load all the necessary bytecode but can then process different programs (e.g., interactively, from within the development environment) without having to be restarted. There exist in any case many solutions to significantly reduce this startup time time (keeping code in memory, optimizing the bytecode reader, reduced versions of CiaoPP that contain only the necessary domains, lazy load, etc.).

**Table 2.** Benchmarks.

<i>boyer</i>	a theorem prover implementation based on Lisp by R. Boyer (nqthm system), performs symbolic evaluation of a given formula;
<i>boyerx</i>	a variant of <i>boyer</i> (using generic term manipulation predicates for formula rewrites);
<i>crypt</i>	cryptomultiplication puzzles solver;
<i>deriv</i>	a program that performs symbolic differentiation of a given formula;
<i>exp</i>	exponential calculation;
<i>factorial</i>	recursive factorial calculation;
<i>fft</i>	fast Fourier transformation calculation;
<i>fib</i>	a program that finds $N$ -th Fibonacci number;
<i>guardians</i>	prison guards game;
<i>hamming</i>	a program that generates the sequence of Hamming numbers;
<i>hanoi</i>	hanoi towers puzzle solver for $N$ disks that are moved over three rods;
<i>jugs</i>	the water jugs problem;
<i>knights</i>	$N$ knights chess problem;
<i>mmatrix</i>	matrix multiplication for two matrices with dimensions $n \times n$ ;
<i>nreverse</i>	naive list reversal;
<i>poly</i>	a program that raises a polynomial $(1 + x + y + z)$ to the 10th power symbolically;
<i>primes</i>	a program that computes $N$ first prime numbers;
<i>progeom</i>	a program that constructs a perfect difference set of order $N$ ;
<i>queens</i>	the $N$ queens program, the number of the queens being the input;
<i>qsort</i>	the quicksort program;
<i>serialize</i>	a palindrome program;
<i>tak</i>	a program that computes the <i>tak</i> function;
<i>witt</i>	the WITT clustering system implementation;

**Table 3.** Benchmark metrics.

Name	Benchmarks		Assertions
	LOC	Size (KB)	total
<i>boyer</i>	853	70	25
<i>boyerx</i>	853	50	23
<i>crypt</i>	76	10	16
<i>deriv</i>	29	8.5	3
<i>exp</i>	28	5.5	5
<i>factorial</i>	13	4.4	3
<i>fft</i>	104	13.2	19
<i>fib</i>	11	4.6	5
<i>guardians</i>	78	9	13
<i>hamming</i>	71	8.7	19
<i>hanoi</i>	44	5.7	4
<i>jugs</i>	132	9.9	9
<i>knights</i>	49	7.5	13
<i>mmatrix</i>	48	5.8	6
<i>nreverse</i>	14	4.8	5
<i>poly</i>	81	12	14
<i>primes</i>	33	6.1	8
<i>progeom</i>	71	8.0	16
<i>qsort</i>	46	6.0	8
<i>queens</i>	47	6.4	10
<i>serialize</i>	81	9.5	9
<i>tak</i>	18	4.6	3
<i>witt</i>	651	50	83

**Table 4.** Benchmarks: full compilation time.

Benchmark	Compilation time, ms			
	Unsafe	Safe		
		Client	RT	CT+RT
<i>boyer</i>	215	1413	1222	393869
<i>boyerx</i>	194	1193	1191	25275
<i>crypt</i>	148	753	755	3776
<i>deriv</i>	142	655	702	993
<i>exp</i>	141	691	684	976
<i>factorial</i>	139	650	658	1207
<i>fft</i>	154	835	827	2757
<i>fib</i>	139	574	664	1154
<i>guardians</i>	146	767	775	1323
<i>hamming</i>	162	796	713	1435
<i>hanoi</i>	142	662	700	834
<i>jugs</i>	142	728	748	1277
<i>knights</i>	153	748	724	1347
<i>mmatrix</i>	146	697	514	982
<i>nreverse</i>	139	662	686	882
<i>poly</i>	147	626	788	1755
<i>primes</i>	139	690	696	998
<i>progeom</i>	140	776	606	1492
<i>qsort</i>	136	712	692	985
<i>queens</i>	149	700	554	1143
<i>serialize</i>	145	735	721	1353
<i>tak</i>	140	604	621	751
<i>witt</i>	219	1418	1558	149855

analysis and assertion checking times. The experiments were run on a MacBook Pro with 2.3GHz Intel Core i7 processor, 16GB RAM, and under the Mac OS X 10.11.1 operating system.

Table 5 shows the detailed analysis times for the *Safe-CT-RT* mode for the benchmarks. The *prep* columns indicate the time needed to load and prepare the analysis, and the *shfr* and *eterms* columns the time to perform sharing and freeness and regular type

**Table 5.** Static analysis time for benchmarks using the *Safe-CT-RT* checking mode (part of total compilation time).

Benchmark	Analysis time, ms				Assertion checking
	prep	shfr	prep	eterms	
boyer	8.185	56.577	8.469	634.059	521.067
boyerx	5.401	42.5	5.456	467.579	343.906
crypt	1.397	8.1	1.359	35.164	115.782
deriv	0.711	2.426	0.669	13.877	24.63
exp	0.402	1.593	0.373	12.471	42.513
factorial	0.276	0.906	0.193	9.3	11.032
fft	1.544	7.425	1.758	36.885	143.532
fib	0.266	1.083	0.215	11.24	13.475
guardians	1.059	4.455	1.049	19.958	46.498
hamming	1.044	6.46	1.022	21.417	60.509
hanoi	0.496	2.267	0.477	11.629	17.453
jugs	0.837	4.039	0.981	22.661	101.27
knights	0.774	3.483	0.748	22.24	46.713
mmatrix	0.48	2.359	0.449	11.671	23.01
nreverse	0.278	2.673	0.253	2.763	7.489
poly	1.414	41.274	1.419	42.363	79.432
primes	0.571	2.278	0.486	12.547	25.364
progeom	0.933	5.6	0.891	21.057	47.704
qsort	0.526	3.725	0.504	6.562	18.491
queens	0.614	3.448	0.587	15.489	35.767
serialize	1.076	11.775	0.984	17.515	46.347
tak	0.327	1.427	0.271	9.986	19.747
witt	13.622	104625.161	14.43	666.404	1079.152

analyses, respectively. The analysis is actually relatively inexpensive compared to the rest of the compilation passes for most of the benchmarks. The regular type analysis is expensive in *boyer* and *boyerx*. The analysis of the formula rewrite predicates generates many large types whose manipulation is expensive. The *witt* benchmark, despite having more regular data structures (tables of sets and matrices), is also expensive to analyze due to a large number of operations. Note that in any case more efficient—but less precise—domains are available for these cases, many within CiaoPP, such as, for example, several widenings for sharing (Navas et al. 2006; Méndez-Lojo et al. 2008), pair sharing domains (Søndergaard 1986; Secci and Spoto 2005), or other type inference domains (Gallagher and de Waal 1994; Bruynooghe and Gallagher 2004).

Finally, Table 6 reports on the execution times for each benchmark in the different assertion checking modes, together with the statistics of assertion checking results. For some of the benchmarks, inputs were varied and this is reflected by the notation *Name(Input)*. The ‘Checked/Total Assrts’ column reports the numbers of statically checked assertions (in the *Safe-CT-RT* checking mode) and total number of assertions for each benchmark. In the worst case the overhead in the *Safe-RT* checking mode is two orders of magnitude higher than in *Client-safe*, but *Safe-CT-RT* removes one order of magnitude (*boyerx*, *fft*, *knights*, *witt*). This is expected since run-time checks of complex properties like data shapes cannot be performed in constant time. The checking changes the complexity of the programs and the overhead increases as the size of the input grows. Note that the *Client-safe* mode also represents the theoretically lowest overhead that we could obtain (assuming a fixed implementation of the instrumentation), by removing internal checks. Sometimes, due to measurement imprecision, the time in *Safe-CT-RT* mode is smaller than in *Client-safe* mode (*crypt*, *exp*, *primes*). In practice, in many programs *Safe-CT-RT* is able to remove most of the checks, except those corresponding to the external predicates. We included in the benchmarks two versions of *boyer*. The original translation (which we call here

*boyerx*) uses *functor/3* and *arg/3* to implement rewrites of arbitrary terms representing formulas. This makes the domains lose precision. The *boyer* version used instead a larger predicate that explicitly enumerates possible formula terms.

## 7. Conclusions

We have addressed the problem of run-time overhead reduction in the context of verification frameworks that combine static and dynamic verification, i.e., systems that combine compile-time and run-time checking of user-provided assertions. Our ultimate objective is to construct an automatic verification system for non-trivial, structural properties, that can be used routinely in production code.

We have defined four practical assertion checking modes, and studied the corresponding trade-offs between the level of guarantees provided by each one and the corresponding execution time slowdown. For these checking modes we have explored the effectiveness of abstract interpretation in detecting the parts of the program’s (partial) specifications that can be statically simplified to true or false, as well as the practical impact of such analysis in reducing the cost of the run-time checks required for the remaining parts of the specifications.

Our experiments have shown that there is indeed a significant advantage in using analysis to reduce run-time tests. We argue that the results are encouraging, supporting the hypothesis that the combination of run-time checking with analysis can reduce checking overhead sufficiently to allow providing full safety in production code, for non-trivial properties. It is worth noticing that the analysis on more complex code like *boyer* allowed us to spot bugs in the original translation from LISP that had been around for 30 years.

We have presented for concreteness our approach in the context of the Ciao language, but the Ciao approach to combining static and dynamic analysis is general and system-independent, as well as the techniques used herein, so we expect the results should carry over to other (dynamic) declarative or imperative languages.

**Table 6.** Benchmark execution times under the different modes and checked vs. total assertions.

Benchmark	Execution time, ms				Checked/Total Assrts	
	Unsafe	Safe			CT+RT	Total
		Client	RT	CT+RT		
boyer	10.404	10.390	2412.197	12.545	25	25
boyerx	15.185	15.511	2284.085	1098.419	21	23
crypt	0.098	0.108	6.196	0.105	15	16
deriv	0.110	0.730	4.640	0.640	2	3
exp	3.838	3.971	63.210	3.899	4	5
factorial	0.008	0.013	0.748	0.013	2	3
fft	23.234	23.972	25879.525	202.655	17	19
fib	0.067	0.073	13.069	0.072	4	5
guardians	2.960	2.980	5805.539	3.221	12	13
hamming	15.259	15.071	7431.143	17.012	18	19
hanoi (2)	0.001	0.011	0.139	0.011		
hanoi (4)	0.002	0.012	1.324	0.013	3	4
hanoi (8)	0.054	0.064	106.649	0.078		
jugs	0.014	0.022	1.660	0.023	8	9
knights	201.278	192.987	17979.597	202.061	12	13
mmatrix (2)	0.001	0.007	0.127	0.008		
mmatrix (3)	0.002	0.010	0.319	0.010	5	6
mmatrix (4)	0.005	0.015	0.662	0.016		
nreverse	2.243	2.313	8217.188	3.125	4	5
poly	1.084	1.272	356.360	1.302	13	14
primes	0.027	0.041	8.860	0.036	7	8
progeom (2)	0.002	0.006	0.630	0.005		
progeom (4)	0.089	0.093	24.424	0.103	15	16
progeom (8)	5.274	5.268	1890.538	5.934		
qsort (8)	0.003	0.006	0.929	0.007		
qsort (16)	0.008	0.013	2.396	0.015	7	8
qsort (32)	0.021	0.028	7.477	0.033		
queens (4)	0.006	0.009	1.076	0.010		
queens (6)	0.122	0.123	20.605	0.139	9	10
queens (8)	2.302	2.281	360.660	2.559		
serialize (9)	0.002	0.007	0.803	0.008		
serialize (16)	0.005	0.011	2.207	0.012	8	9
serialize (25)	0.010	0.017	5.029	0.019		
tak	2.855	2.866	963.481	3.665	2	3
witt	15.886	15.947	1511.059	316.063	69	83

## References

- N. Bjørner, F. Fioravanti, A. Rybalchenko, and V. Senni, editors. *Workshop on Horn Clauses for Verification and Synthesis*, July 2014. URL <http://vs12014.at/meetings/HCVS-index.html>. To appear in Electronic Proceedings in Theoretical Computer Science.
- J. Boye, W. Drabent, and J. Maluszyński. Declarative Diagnosis of Constraint Programs: an assertion-based approach. In *Proc. of the 3rd. Int'l Workshop on Automated Debugging-AADEBUD'97*, pages 123–141, Linköping, Sweden, May 1997. U. of Linköping Press.
- M. Bruynooghe and J. Gallagher. Inferring Polymorphic Types from Logic Programs. In *International Symposium on Logic-based Program Synthesis and Transformation (LOPSTR'04)*. Preproceedings, July 2004.
- F. Bueno, P. Deransart, W. Drabent, G. Ferrand, M. Hermenegildo, J. Maluszyński, and G. Puebla. On the Role of Semantic Approximations in Validation and Diagnosis of Constraint Logic Programs. In *Proc. of the 3rd. Int'l WS on Automated Debugging-AADEBUD*, pages 155–170. U. Linköping Press, May 1997.
- R. Cartwright and M. Fagan. Soft Typing. In *PLDI'91*, pages 278–292. SIGPLAN, ACM, 1991.
- P. Cousot and R. Cousot. Abstract Interpretation: a Unified Lattice Model for Static Analysis of Programs by Construction or Approximation of Fixpoints. In *ACM Symposium on Principles of Programming Languages (POPL'77)*. ACM Press, 1977.
- C. Dimoulas and M. Felleisen. On contract satisfaction in a higher-order world. *ACM Trans. Program. Lang. Syst.*, 33(5):16, 2011.
- W. Drabent, S. Nadjm-Tehrani, and J. Maluszyński. The Use of Assertions in Algorithmic Debugging. In *Intl. Conf. on Fifth Generation Computer Systems*, pages 573–581, 1988.
- M. Fähndrich and F. Logozzo. Static contract checking with abstract interpretation. In *Proceedings of the 2010 International Conference on Formal Verification of Object-oriented Software*, volume 6528 of *FoVeOOS'10*, pages 10–30, Berlin, Heidelberg, 2011. Springer-Verlag. ISBN 3-642-18069-8, 978-3-642-18069-9. URL <http://dl.acm.org/citation.cfm?id=1949303.1949305>.
- R. B. Findler and M. Felleisen. Contracts for higher-order functions. In M. Wand and S. L. P. Jones, editors, *ICFP*, pages 48–59. ACM, 2002. ISBN 1-58113-487-8.
- J. Gallagher and D. de Waal. Fast and precise regular approximations of logic programs. In P. Van Hentenryck, editor, *Proc. of the 11th International Conference on Logic Programming*, pages 599–613. MIT Press, 1994.
- S. Grebenshchikov, A. Gupta, N. P. Lopes, C. Popeea, and A. Rybalchenko. HSF(C): A Software Verifier Based on Horn Clauses - (Competition Contribution). In C. Flanagan and B. König, editors, *TACAS*, volume 7214 of *LNCS*, pages 549–551. Springer, 2012. ISBN 978-3-642-28755-8.

- A. Gurfinkel, T. Kahsai, A. Komuravelli, and J. A. Navas. The seahorn verification framework. In D. Kroening and C. S. Pasareanu, editors, *Computer Aided Verification - 27th International Conference, CAV 2015, San Francisco, CA, USA, July 18-24, 2015, Proceedings, Part I*, volume 9206 of *Lecture Notes in Computer Science*, pages 343–361. Springer, 2015. ISBN 978-3-319-21689-8. doi: 10.1007/978-3-319-21690-4\_20. URL [http://dx.doi.org/10.1007/978-3-319-21690-4\\_20](http://dx.doi.org/10.1007/978-3-319-21690-4_20).
- M. Hermenegildo, G. Puebla, and F. Bueno. Using Global Analysis, Partial Specifications, and an Extensible Assertion Language for Program Validation and Debugging. In K. R. Apt, V. Marek, M. Truszczynski, and D. S. Warren, editors, *The Logic Programming Paradigm: a 25-Year Perspective*, pages 161–192. Springer-Verlag, July 1999.
- M. Hermenegildo, G. Puebla, F. Bueno, and P. L. García. Integrated Program Debugging, Verification, and Optimization Using Abstract Interpretation (and The Ciao System Preprocessor). *Science of Computer Programming*, 58(1–2):115–140, 2005.
- M. V. Hermenegildo, F. Bueno, M. Carro, P. López, E. Mera, J. Morales, and G. Puebla. The Ciao Approach to the Dynamic vs. Static Language Dilemma. In *Proceedings for the International Workshop on Scripts to Programs, STOP'11*, New York, NY, USA, 2011. ACM.
- M. V. Hermenegildo, F. Bueno, M. Carro, P. López, E. Mera, J. Morales, and G. Puebla. An Overview of Ciao and its Design Philosophy. *Theory and Practice of Logic Programming*, 12(1–2):219–252, January 2012. doi: doi:10.1017/S1471068411000457. <http://arxiv.org/abs/1102.5497>.
- E. Koukoutos and V. Kuncak. Checking Data Structure Properties Orders of Magnitude Faster. In B. Bonakdarpour and S. A. Smolka, editors, *Runtime Verification*, volume 8734 of *Lecture Notes in Computer Science*, pages 263–268. Springer International Publishing, 2014. ISBN 978-3-319-11163-6. doi: 10.1007/978-3-319-11164-3\_22. URL [http://dx.doi.org/10.1007/978-3-319-11164-3\\_22](http://dx.doi.org/10.1007/978-3-319-11164-3_22).
- C. Lai. Assertions with Constraints for CLP Debugging. In P. Deransart, M. V. Hermenegildo, and J. Maluszynski, editors, *Analysis and Visualization Tools for Constraint Programming*, volume 1870 of *Lecture Notes in Computer Science*, pages 109–120. Springer, 2000. ISBN 3-540-41137-2.
- L. Lamport and L. C. Paulson. Should your specification language be typed? *ACM Transactions on Programming Languages and Systems*, 21(3):502–526, May 1999.
- G. T. Leavens, K. R. M. Leino, and P. Müller. Specification and verification challenges for sequential object-oriented programs. *Formal Asp. Comput.*, 19(2):159–189, 2007.
- U. Liqat, S. Kerrison, A. Serrano, K. Georgiou, P. Lopez-Garcia, N. Grech, M. Hermenegildo, and K. Eder. Energy Consumption Analysis of Programs based on X MOS ISA-level Models. In G. Gupta and R. Peña, editors, *Logic-Based Program Synthesis and Transformation, 23rd International Symposium, LOPSTR 2013, Revised Selected Papers*, volume 8901 of *Lecture Notes in Computer Science*, pages 72–90. Springer, 2014. ISBN 978-3-319-14124-4. doi: 10.1007/978-3-319-14125-1\_5. URL [http://dx.doi.org/10.1007/978-3-319-14125-1\\_5](http://dx.doi.org/10.1007/978-3-319-14125-1_5).
- U. Liqat, K. Georgiou, S. Kerrison, P. Lopez-Garcia, M. V. Hermenegildo, J. P. Gallagher, and K. Eder. Inferring Parametric Energy Consumption Functions at Different Software Levels: ISA vs. LLVM IR. In M. V. Eekelen and U. D. Lago, editors, *Foundational and Practical Aspects of Resource Analysis. Fourth International Workshop FOPARA 2015, Revised Selected Papers*, Lecture Notes in Computer Science. Springer, 2016. URL <http://arxiv.org/abs/1511.01413>. In press.
- M. Méndez-Lojo, J. Navas, and M. Hermenegildo. A Flexible (C)LP-Based Approach to the Analysis of Object-Oriented Programs. In *17th International Symposium on Logic-based Program Synthesis and Transformation (LOPSTR 2007)*, number 4915 in *Lecture Notes in Computer Science*, pages 154–168. Springer-Verlag, August 2007.
- M. Méndez-Lojo, O. Lhoták, and M. Hermenegildo. Efficient Set Sharing using ZBDDs. In *21st Int'l. WS on Languages and Compilers for Parallel Computing (LCPC'08)*, LNCS. Springer-Verlag, August 2008.
- E. Mera, P. López-García, and M. Hermenegildo. Integrating Software Testing and Run-Time Checking in an Assertion Verification Framework. In *25th Int'l. Conference on Logic Programming (ICLP'09)*, number 5649 in LNCS, pages 281–295. Springer-Verlag, July 2009.
- E. Mera, T. Trigo, P. López-García, and M. Hermenegildo. Profiling for Run-Time Checking of Computational Properties and Performance Debugging. In *Practical Aspects of Declarative Languages (PADL'11)*, volume 6539 of *Lecture Notes in Computer Science*, pages 38–53. Springer-Verlag, January 2011.
- K. Muthukumar and M. Hermenegildo. Determination of Variable Dependence Information at Compile-Time Through Abstract Interpretation. In *1989 North American Conference on Logic Programming*, pages 166–189. MIT Press, October 1989.
- K. Muthukumar and M. Hermenegildo. Deriving A Fixpoint Computation Algorithm for Top-down Abstract Interpretation of Logic Programs. Technical Report ACT-DC-153-90, Microelectronics and Computer Technology Corporation (MCC), Austin, TX 78759, April 1990. URL <ftp://cliplab.org/pub/papers/tr153-90.mcc.ps.Z>.
- K. Muthukumar and M. Hermenegildo. Combined Determination of Sharing and Freeness of Program Variables Through Abstract Interpretation. In *International Conference on Logic Programming (ICLP 1991)*, pages 49–63. MIT Press, June 1991.
- K. Muthukumar and M. Hermenegildo. Compile-time Derivation of Variable Dependency Using Abstract Interpretation. *Journal of Logic Programming*, 13(2/3):315–347, July 1992.
- J. Navas, F. Bueno, and M. Hermenegildo. Efficient top-down set-sharing analysis using cliques. In *Eight International Symposium on Practical Aspects of Declarative Languages*, number 2819 in LNCS, pages 183–198. Springer-Verlag, January 2006.
- J. Navas, M. Méndez-Lojo, and M. Hermenegildo. Safe Upper-bounds Inference of Energy Consumption for Java Bytecode Applications. In *The Sixth NASA Langley Formal Methods Workshop (LFM 08)*, pages 29–32, April 2008. Extended Abstract.
- J. Navas, M. Méndez-Lojo, and M. Hermenegildo. User-Definable Resource Usage Bounds Analysis for Java Bytecode. In *Proceedings of the Workshop on Bytecode Semantics, Verification, Analysis and Transformation (BYTECODE'09)*, volume 253 of *Electronic Notes in Theoretical Computer Science*, pages 65–82. Elsevier - North Holland, March 2009.
- P. Pietrzak, J. Correia, G. Puebla, and M. Hermenegildo. Context-Sensitive Multivariant Assertion Checking in Modular Programs. In *13th International Conference on Logic for Programming Artificial Intelligence and Reasoning (LPAR'06)*, number 4246 in LNCS, pages 392–406. Springer-Verlag, November 2006.
- G. Puebla, F. Bueno, and M. Hermenegildo. An Assertion Language for Debugging of Constraint Logic Programs. In *ILPS'97 WS on Tools and Environments for (C)LP*, October 1997. [ftp://cliplab.org/pub/papers/assert\\_lang\\_tr\\_discipldeliv.ps.gz](ftp://cliplab.org/pub/papers/assert_lang_tr_discipldeliv.ps.gz).
- G. Puebla, F. Bueno, and M. Hermenegildo. An Assertion Language for Constraint Logic Programs. In P. Deransart, M. Hermenegildo, and J. Maluszynski, editors, *Analysis and Visualization Tools for Constraint Programming*, number 1870 in LNCS, pages 23–61. Springer-Verlag, September 2000a.
- G. Puebla, F. Bueno, and M. Hermenegildo. Combined Static and Dynamic Assertion-Based Debugging of Constraint Logic Programs. In *Logic-based Program Synthesis and Transformation (LOPSTR'99)*, number 1817 in LNCS, pages 273–292. Springer-Verlag, March 2000b.
- A. Rastogi, N. Swamy, C. Fournet, G. M. Bierman, and P. Vekris. Safe & efficient gradual typing for typescript. In S. K. Rajamani and D. Walker, editors, *Proceedings of the 42nd Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, POPL 2015, Mumbai, India, January 15-17, 2015*, pages 167–180. ACM, 2015. ISBN 978-1-4503-3300-9. doi: 10.1145/2676726.2676971. URL <http://doi.acm.org/10.1145/2676726.2676971>.
- T. W. Schiller, K. Donohue, F. Coward, and M. D. Ernst. Case studies and tools for contract specifications. In *Proceedings of the 36th International Conference on Software Engineering, ICSE 2014*, pages 596–607. New York, NY, USA, 2014. ACM. ISBN 978-1-4503-2756-5. doi: 10.1145/2568225.2568285. URL <http://doi.acm.org/10.1145/2568225.2568285>.

- S. Secci and F. Spoto. Pair-Sharing Analysis of Object-Oriented Programs. In *12th International Symposium Static Analysis Symposium (SAS'05)*, volume 3672 of *Lecture Notes in Computer Science*. Springer, 2005.
- H. Søndergaard. An application of abstract interpretation of logic programs: occur check reduction. In *European Symposium on Programming, LNCS 123*, pages 327–338. Springer-Verlag, 1986.
- N. Stulova, J. F. Morales, and M. V. Hermenegildo. Assertion-based Debugging of Higher-Order (C)LP Programs. In *16th Int'l. ACM SIGPLAN Symposium on Principles and Practice of Declarative Programming (PPDP'14)*. ACM Press, September 2014.
- N. Stulova, J. F. Morales, and M. V. Hermenegildo. Practical Run-time Checking via Unobtrusive Property Caching. *Theory and Practice of Logic Programming, 31st Int'l. Conference on Logic Programming (ICLP'15) Special Issue*, 15(04-05):726–741, September 2015. ISSN 1471-0684. URL <http://arxiv.org/abs/1507.05986>.
- A. Takikawa, D. Feltey, E. Dean, M. Flatt, R. B. Findler, S. Tobin-Hochstadt, and M. Felleisen. Towards practical gradual typing. In J. T. Boyland, editor, *29th European Conference on Object-Oriented Programming, ECOOP 2015, July 5-10, 2015, Prague, Czech Republic*, volume 37 of *LIPICs*, pages 4–27. Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik, 2015. ISBN 978-3-939897-86-6. doi: 10.4230/LIPICs.ECOOP.2015.4. URL <http://dx.doi.org/10.4230/LIPICs.ECOOP.2015.4>.
- A. Takikawa, D. Feltey, B. Greenman, M. S. New, J. Vitek, and M. Felleisen. Is sound gradual typing dead? In R. Bodík and R. Majumdar, editors, *Proceedings of the 43rd Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, POPL 2016, St. Petersburg, FL, USA, January 20 - 22, 2016*, pages 456–468. ACM, 2016. ISBN 978-1-4503-3549-2. doi: 10.1145/2837614.2837630. URL <http://doi.acm.org/10.1145/2837614.2837630>.
- S. Tobin-Hochstadt and M. Felleisen. The Design and Implementation of Typed Scheme. In *POPL*, pages 395–406. ACM, 2008.
- C. Vaucheret and F. Bueno. More Precise yet Efficient Type Inference for Logic Programs. In *International Static Analysis Symposium*, volume 2477 of *Lecture Notes in Computer Science*, pages 102–116. Springer-Verlag, September 2002.