Generation of Basic Semi-algebraic Invariants Using Convex Polyhedra*

Roberto Bagnara[1], Enric Rodríguez-Carbonell[2], and Enea Zaffanella[1]

[1] Department of Mathematics, University of Parma, Italy
{bagnara,zaffanella}@cs.unipr.it
[2] Software Department, Technical University of Catalonia, Spain
erodri@lsi.upc.edu

Abstract. A technique for generating invariant polynomial *inequalities* of bounded degree is presented using the abstract interpretation framework. It is based on overapproximating basic semi-algebraic sets, i.e., sets defined by conjunctions of polynomial inequalities, by means of convex polyhedra. While improving on the existing methods for generating invariant polynomial *equalities*, since polynomial inequalities are allowed in the guards of the transition system, the approach does not suffer from the prohibitive complexity of methods based on quantifier-elimination. The application of our implementation to benchmark programs shows that the method produces non-trivial invariants in reasonable time. In some cases the generated invariants are essential to verify safety properties that cannot be proved with just classical linear invariants.

1 Introduction

The discovery of invariant properties is at the core of the analysis and verification of infinite state systems such as sequential programs and reactive systems. For this reason invariant generation has been a major research problem since the seventies. Abstract interpretation [12] provides a solid foundation for the development of techniques automatizing the synthesis of invariants of several classes, most significantly intervals [11], linear equalities [25] and linear inequalities [15].

For some applications, linear invariants are not enough to get a precise analysis of numerical programs and non-linear invariants may be needed as well. For example, the ASTRÉE static analyzer, which has been successfully employed to verify the absence of run-time errors in flight control software [14], implements the ellipsoid abstract domain [7], which represents a certain class of quadratic inequality invariants. Moreover, it has been acknowledged elsewhere [33, 35] that non-linear invariants are sometimes required to prove program properties.

As a consequence, a remarkable amount of work has been recently directed to the generation of invariant *polynomial equalities*. Some of the methods plainly

* This work has been partially supported by PRIN project “AIDA — Abstract Interpretation: Design and Applications,” by the “LogicTools” project (CICYT TIN 2004-03382), and the FPU grant AP2002-3693 from the Spanish MECD.

ignore all the conditional guards [30, 32]; other methods can only consider the polynomial equalities in the guards [9, 36], whereas some other proposals [28, 31] can handle *polynomial disequalities* in guards (i.e., guards of the form $p \neq 0$ where p is a polynomial). All the techniques previously mentioned cannot handle the case of *polynomial inequalities* in the guards: these are ignored to the expense of precision.

In this paper we present a method for generating conjunctions of polynomial inequalities as invariants of transition systems, which we have chosen as our programming model. The transition systems that the approach can handle admit conjunctions of polynomial inequalities as guards and initial conditions, as well as polynomial assignments and nondeterministic assignments where the *rvalue* is unknown (which may correspond, for instance, to the assignment of expressions that cannot be modeled by means of polynomials).

Formally, our technique is an abstract interpretation in the lattice of *polynomial cones* of bounded degree, which are the algebraic structures analogous to vector spaces in the context of polynomial equality invariants [9]. Intuitively, the approach is based on considering nonlinear terms as additional independent variables, and then work with convex polyhedra, which are used to represent polynomial cones, in this extended set of variables. In order to reduce the loss of precision induced by this overapproximation, additional linear constraints, relating the newly added variables with the original ones, are added conservatively to the polyhedra, so as to enforce some (semantically redundant) nonlinear constraints that would be lost in the translation. The strength of the approach is that, while allowing for a much broader class of programs than linear analysis, it uses the very same underlying machinery: this permits the adoption of already existing implementations of convex polyhedra like [4], as well as the possibility of resorting to further approximations, such as *bounded differences* [1] or *octagons* [27], when facing serious scalability problems.

The rest of the paper is organized as follows. In the next subsection, related work is briefly reviewed. Section 2 gives background information on algebraic geometry, transition systems and abstract interpretation. In Section 3 we argue why the first approach we propose, based on abstract interpretation over the first-order language of polynomial inequalities, is not feasible. Section 4 presents the main contribution of the paper, where it is shown how polynomial inequalities can be discovered as invariants by means of polynomial cones, represented as convex polyhedra. The experimental evaluation of our implementation of these ideas is described in Section 5. Finally in Section 6 we summarize the contributions of the paper and sketch some ideas for future work.

1.1 Related Work

To the best of our knowledge, the first contribution towards the generation of invariant polynomial inequalities is [6]. The authors consider a simple class of transition systems, where assignments are of the form $x := x + k$ or $x := k$ with $k \in \mathbb{Z}$. Such a transition system is soundly abstracted into a new one whose exact reachability set is computable and overapproximates the reachability set of the original system. Besides the fact that the programming model is more restrictive than the one used in this paper, these ideas do not seem to have undergone experimental evaluation so that, as far as we can tell, their practical value remains to be assessed.

In [24], Kapur proposes a method based on imposing that a template polynomial inequality with undetermined coefficients is invariant and solving the resulting constraints over the coefficients by real quantifier elimination. Unfortunately, the great computational complexity of quantifier elimination appears to make the method impractical: as the author reports, an experimental implementation performed poorly or did not return any answer for all the analyzed programs [D. Kapur, personal communication, 2005].

A similar idea is at the core of [10, 33], where, instead of real quantifier elimination, semidefinite programming is employed. The method, which is reported to perform rather efficiently for several interesting cases, automatically determines *one* solution to the constraint system on the template parameters. This is particularly appropriate for proving program termination because, once a class of candidate ranking functions has been chosen, any solution belonging to this class is good enough. The same approach has also been applied to the computation of invariant properties. In this case, according to [10], the one above becomes the main limitation of the method: any invariant property, even a weak one, may be obtained and it is unclear whether it is possible to drive the solver so as to produce a more precise invariant in the same class.

In [35], Sankaranarayanan et al. propose a technique for generating linear invariants by linear programming. It is based on imposing, as invariants, constraints where the coefficients of the variables are fixed *a priori*; the analysis then returns, for each such constraint, an independent term for which the constraint is indeed an invariant of the system (in the case where this is not possible, the analysis returns $\pm\infty$). A generalization of this approach for the discovery of invariant polynomial inequalities by means of semidefinite programming is sketched. Similarly, in [7] the ellipsoid abstract domain is presented, which allows to generate invariant quadratic inequalities with two variables by also fixing the coefficients of terms and leaving the independent term to be determined by the analysis. The approach proposed in this paper differs in that we do not need to fix any of these coefficients in advance, but rather it is the analysis itself that determines all coefficients.

2 Preliminaries

2.1 Algebraic Geometry

We denote the real numbers by $\mathbb{R}$, and the nonnegative real numbers by $\mathbb{R}_+$. A *term* in the tuple of variables $\boldsymbol{x} = (x_1, \ldots, x_n)$ is an expression of the form $\boldsymbol{x}^{\boldsymbol{\alpha}} = x_1^{\alpha_1} x_2^{\alpha_2} \cdots x_n^{\alpha_n}$, where $\boldsymbol{\alpha} = (\alpha_1, \ldots, \alpha_n) \in \mathbb{N}^n$. A *monomial* is an expression of the form $c \cdot \boldsymbol{x}^{\boldsymbol{\alpha}}$, simply written as $c\boldsymbol{x}^{\boldsymbol{\alpha}}$, where $c \in \mathbb{R}$ and $\boldsymbol{x}^{\boldsymbol{\alpha}}$ is a term. The *degree* of a monomial $c\boldsymbol{x}^{\boldsymbol{\alpha}}$ with $c \neq 0$ is $\deg(c\boldsymbol{x}^{\boldsymbol{\alpha}}) := \alpha_1 + \cdots + \alpha_n$; the degree of 0 is $\deg(0) := -\infty$. A *polynomial* is a finite sum of monomials. The set of all polynomials in $\boldsymbol{x}$ with coefficients in $\mathbb{R}$ is denoted by $\mathbb{R}[\boldsymbol{x}]$. The degree of a non-null polynomial is the maximum of the degrees of its monomials. We denote by $\mathbb{R}_d[\boldsymbol{x}]$ the set of all polynomials in $\mathbb{R}[\boldsymbol{x}]$ having degree at most d. In particular, the polynomials in $\mathbb{R}_1[\boldsymbol{x}]$, i.e., having degree at most 1, are called *linear*; similarly, the polynomials in $\mathbb{R}_2[\boldsymbol{x}]$ are called *quadratic*.

A *polynomial equality* (resp., *polynomial inequality*) is a formula of the form $p = 0$ (resp., $p \geq 0$), where $p \in \mathbb{R}[\boldsymbol{x}]$. Both will be referred to as *polynomial constraints* or simply *constraints*. Given a constraint system ψ, i.e., a finite set of polynomial constraints, we define

$$\operatorname{poly}(\psi) := \bigl\{\, p \in \mathbb{R}[\boldsymbol{x}] \bigm| (p = 0) \in \psi \text{ or } (-p = 0) \in \psi \text{ or } (p \geq 0) \in \psi \,\bigr\}.$$

We will sometimes abuse notation by writing the set ψ to denote the finite conjunction of the constraints occurring in it.

The *algebraic set* defined by a finite set of polynomials $\{p_1, \ldots, p_k\} \subseteq \mathbb{R}[\boldsymbol{x}]$ is the set of points that satisfy the corresponding polynomial equalities, i.e.,

$$\bigl\{\, \boldsymbol{v} \in \mathbb{R}^n \bigm| p_1(\boldsymbol{v}) = 0, \ldots, p_k(\boldsymbol{v}) = 0 \,\bigr\}.$$

Similarly, the *basic semi-algebraic set* defined by the same set of polynomials is the set of points that satisfy all the corresponding polynomial inequalities:

$$\bigl\{\, \boldsymbol{v} \in \mathbb{R}^n \bigm| p_1(\boldsymbol{v}) \geq 0, \ldots, p_k(\boldsymbol{v}) \geq 0 \,\bigr\}.$$

Finally, *semi-algebraic sets* are obtained from basic semi-algebraic sets by taking complements, finite unions and finite intersections.

2.2 Transition Systems

In this section we define our programming model: *transition systems*.

Definition 1. (Transition system.) *A* transition system $(\boldsymbol{x}, \mathcal{L}, \mathcal{T}, \mathcal{I})$ *is a tuple that consists of the following components:*

- *An n-tuple of real-valued* variables $\boldsymbol{x} = (x_1, \ldots, x_n)$.
- *A finite set $\mathcal{L}$ of* locations.
- *A finite set $\mathcal{T} \subset \mathcal{L} \times \mathcal{L} \times \wp(\mathbb{R}^n) \times \bigl(\mathbb{R}^n \to \wp(\mathbb{R}^n)\bigr)$ of* transitions. *A transition $(\ell, \ell', \gamma, \rho) \in \mathcal{T}$ consists of a* source location $\ell \in \mathcal{L}$, *a* target location $\ell' \in \mathcal{L}$, *a* guard $\gamma \subseteq \mathbb{R}^n$ *that enables the transition, and, finally, an* update map $\rho \colon \mathbb{R}^n \to \wp(\mathbb{R}^n)$ *that relates the values of the variables before and after the firing of the transition.*
- *A map $\mathcal{I} \colon \mathcal{L} \to \wp(\mathbb{R}^n)$ from locations to* initial conditions.

The guards, the update maps and the initial conditions are all assumed to be finitely computable.

The state of a transition system is completely characterized by the location at which control resides and by a valuation for the variables.

Definition 2. (Local and global state.) *A* local state *(at some unspecified location) is any real vector $\boldsymbol{v} = (v_1, \dots, v_n) \in \mathbb{R}^n$, interpreted as a valuation for the variables $\boldsymbol{x} = (x_1, \dots, x_n)$: in local state $\boldsymbol{v}$, we have $x_i = v_i$ for each $i = 1$, …, n. A* global state *is a pair $(\ell, \boldsymbol{v})$, where $\ell \in \mathcal{L}$ and $\boldsymbol{v}$ is the local state at ℓ.*

Definition 3. (Run, initial state.) *A* run *of the transition system $(\boldsymbol{x}, \mathcal{L}, \mathcal{T}, \mathcal{I})$ is a sequence of global states $(\ell_0, \boldsymbol{v}_0)$, $(\ell_1, \boldsymbol{v}_1)$, $(\ell_2, \boldsymbol{v}_2)$, … such that (1) $(\ell_0, \boldsymbol{v}_0)$ is an* initial state*, that is $\boldsymbol{v}_0 \in \mathcal{I}(\ell_0)$, and (2) for each pair of consecutive states, $(\ell_i, \boldsymbol{v}_i)$ and $(\ell_{i+1}, \boldsymbol{v}_{i+1})$, there exists a transition $(\ell_i, \ell_{i+1}, \gamma, \rho) \in \mathcal{T}$ that is enabled, i.e., $\boldsymbol{v}_i \in \gamma$, and such that $\boldsymbol{v}_{i+1} \in \rho(\boldsymbol{v}_i)$.*

The fundamental notion is that of an *invariant* of a transition system:

Definition 4. (Reachable state, invariant property and map.) *A global state $(\ell, \boldsymbol{v})$ is called* reachable *in the transition system $S = (\boldsymbol{x}, \mathcal{L}, \mathcal{T}, \mathcal{I})$, if there exists a run $(\ell_0, \boldsymbol{v}_0)$, $(\ell_1, \boldsymbol{v}_1)$, …, $(\ell_m, \boldsymbol{v}_m)$ of S such that $(\ell, \boldsymbol{v}) = (\ell_m, \boldsymbol{v}_m)$. We denote the set of reachable states of S by $\mathrm{reach}(S)$, and the set of (local) reachable states at location ℓ, i.e., those $\boldsymbol{v}$ such that $(\ell, \boldsymbol{v}) \in \mathrm{reach}(S)$, by $\mathrm{reach}_\ell(S)$.*

If $\boldsymbol{x} = (x_1, \dots, x_n)$, an invariant property *of S at location $\ell \in \mathcal{L}$ (also called an* invariant*) is any set $I \in \wp(\mathbb{R}^n)$ such that $\mathrm{reach}_\ell(S) \subseteq I$. Finally, an* invariant map *is a map $\mathrm{inv} \colon \mathcal{L} \to \wp(\mathbb{R}^n)$ such that for any $\ell \in \mathcal{L}$, $\mathrm{inv}(\ell)$ is an invariant of S at location ℓ.*

In this paper we focus on a particular class of transition systems, *basic semi-algebraic transition systems*:

Definition 5. (Basic semi-algebraic transition system.) *A transition system $(\boldsymbol{x}, \mathcal{L}, \mathcal{T}, \mathcal{I})$, where $\boldsymbol{x} = (x_1, \dots, x_n)$, is called* basic semi-algebraic *if:*

1. *for all $(\ell, \ell', \gamma, \rho) \in \mathcal{T}$, γ is a basic semi-algebraic set and there exist $k \leq n$ polynomials $p_1, \dots, p_k \in \mathbb{R}[\boldsymbol{x}]$ and distinct indices $i_1, \dots, i_k \in \{1, \dots, n\}$ such that, for each $\boldsymbol{v} \in \mathbb{R}^n$,*

$$\rho(\boldsymbol{v}) = \left\{\, (v_1', \dots, v_n') \in \mathbb{R}^n \;\middle|\; v_{i_1}' = p_1(\boldsymbol{v}), \dots, v_{i_k}' = p_k(\boldsymbol{v}) \,\right\};$$

2. *$\mathcal{I}(\ell)$ is a basic semi-algebraic set, for each $\ell \in \mathcal{L}$.*

Notice that a basic semi-algebraic transition system can also model *nondeterministic assignments*, that is, assignments whose *rvalue* is unknown.

Example 1. The program shown on the left of Figure 1 is a minor variant of the program in [18, p. 64], computing the floor of the square root of a natural number a. The basic semi-algebraic transition system shown on the right of the figure models the (second loop of the) program. Note that even the original program in [18], which has the disequality $c \neq 1$ in the loop guard, can be modeled as a basic semi-algebraic transition system (by translating $c \neq 1$ as $c \leq 0 \lor c \geq 2$ and then having four transitions instead of two). The variant in Figure 1 has been adopted just for presentation purposes: its analysis leads to the same invariants that are computed when analyzing the original program.

```
     { $a \geq 0$ }

     $b := 0$; $c := 1$;

     while $c^2 \leq a$ do
         $c := 2c$;
     end while

$\ell_0$ :  while $c \geq 2$ do
         $c := c/2$;
         if $(b+c)^2 \leq a$ then
             $b := b + c$;
         end if
     end while

     return $b$;
```

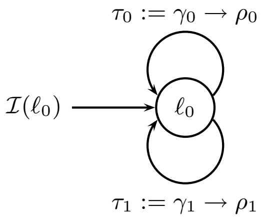
$\tau_0 := \gamma_0 \to \rho_0$
$\mathcal{I}(\ell_0)$
ℓ_0
$\tau_1 := \gamma_1 \to \rho_1$

$$\mathcal{I}(\ell_0) := \{a \geq 0, b = 0, c \geq 1, c^2 \geq a + 1\}$$

$$\gamma_0 := \bigl\{c \geq 2, (2b + c)^2 \leq 4a\bigr\}$$
$$\rho_0 := (a', b', c') = (a, b + c/2, c/2)$$

$$\gamma_1 := \bigl\{c \geq 2, (2b + c)^2 \geq 4a + 4\bigr\}$$
$$\rho_1 := (a', b', c') = (a, b, c/2)$$

Fig. 1. A program and its model as a basic semi-algebraic transition system

2.3 Abstract Interpretation

Abstract interpretation [12] is a general theory of approximation of the behavior of dynamic discrete systems. One of its classical applications is the inference of invariant properties of transition systems [13]. This is done by specifying the set of reachable states of the given transition system as the solution of a system of fixpoint equations. The concrete behavior of the transition system is then overapproximated by setting up a corresponding system of equations defined over an *abstract domain*, providing computable representations for the *abstract properties* that are of interest for the analysis, as well as *abstract operations* that are sound approximations of the *concrete operations* used by the transition system being analyzed. The solution of the system of abstract equations can be found iteratively, possibly applying further conservative approximations and using convergence acceleration methods, such as *widenings* [12]. One of the main advantages of this methodology for the inference of invariant properties is that the correctness of the results thus obtained follows by design.

Given a transition system $S = (\boldsymbol{x}, \mathcal{L}, \mathcal{T}, \mathcal{I})$, the set of its reachable states $\mathrm{reach}(S)$ can be characterized by means of a system of fixpoint equations where, for each $\ell \in \mathcal{L}$, we have the equation

$$\mathrm{reach}_\ell(S) = \mathcal{I}(\ell) \cup \bigcup \Bigl\{ \rho\bigl(\mathrm{reach}_{\ell'}(S) \cap \gamma\bigr) \Bigm| (\ell', \ell, \gamma, \rho) \in \mathcal{T} \Bigr\}. \tag{1}$$

The least fixpoint of this system of equations, with respect to the pointwise extension of the subset ordering on $\wp(\mathbb{R}^n)$, is $\mathrm{reach}(S)$; any overapproximation of $\mathrm{reach}(S)$ yields an invariant map for S. Abstract interpretation provides us with a methodology to constructively obtain an invariant map.

The application of abstract interpretation in this setting involves:

Choosing an abstract domain A. Each element in the abstract domain overapproximates a set of local states. The original system of fixpoint equations $\bar{Y} = \bar{F}(\bar{Y})$, expressing the semantics of the transition system, is transformed into a fixpoint equation $\bar{Z} = \bar{G}(\bar{Z})$ over the abstract values.[3] A convenient (though, in general, not the most precise) way of obtaining $\bar{G}$ from $\bar{F}$ is to follow the syntactic structure of Equation (1): the sets of local states $\mathcal{I}(\ell)$ and γ are mapped into corresponding abstract values, and the concrete operations of *union*, *intersection* and *update*, working on sets of local states, are translated into corresponding operations on the abstract domain. In both cases, the relationship between the concrete and the abstract domains, established according to the theory of abstract interpretation, ensures that the translation is correct and provides an overapproximation of the concrete behavior of the system.

Computing iteratively a post-fixpoint of $\bar{Z} = \bar{G}(\bar{Z})$. Any post-fixpoint of the recursive equation $\bar{Z} = \bar{G}(\bar{Z})$ describes an invariant map for the transition system. The least of these post-fixpoints can be obtained by computing the increasing sequence defined by $\bar{Z}_0 = \bar{\bot}$ (where $\bot$ is the least element of the abstract domain, which usually describes the empty set of states) and, for $k \in \mathbb{N}$, $\bar{Z}_{k+1} = \bar{G}(\bar{Z}_k)$. However, this iterative process may not converge in finite time: termination can be forced by the adoption of a *widening* operator $\nabla \colon A \times A \to A$, at the cost of further overapproximation. If $\sqsubseteq \, \subseteq A \times A$ and $\sqcup \colon A \times A \to A$ are the abstract operators approximating the subset inclusion partial order and the union of sets of states, respectively, then a widening ∇ over A must satisfy the following conditions:

1. for each $a_1, a_2 \in A$ such that $a_1 \sqsubseteq a_2$ we have $a_2 \sqsubseteq a_1 \nabla a_2$;
2. for any increasing chain $a_0 \sqsubseteq a_1 \sqsubseteq \cdots$, the new increasing chain defined by $a'_0 = a_0$, $a'_{k+1} = a'_k \nabla (a'_k \sqcup a_{k+1})$ is not strictly increasing (that is, it finitely converges).

By construction, the abstract iteration sequence with widening converges, in a finite number of steps, to a post-fixpoint of the equation. Notice that the termination test for this analysis method requires the evaluation of the *abstract inclusion* relation '$\sqsubseteq$': iteration can be stopped at the smallest $i \in \mathbb{N}$ such that $\bar{Z}'_{i+1} \sqsubseteq \bar{Z}'_i$.

3 A Possible Approach with Real Quantifier Elimination

A first approach to generate polynomial inequalities as invariants would be to use semi-algebraic sets as abstract values, i.e., as overapproximations of the states of the system. It is a well-known result by Tarski [38] that the first-order theory of polynomial inequalities over the reals admits computable quantifier elimination: that is, given a formula $Q_1 y_1 \ldots . Q_m y_m \, . \, \psi(x_1, \ldots, x_n, y_1, \ldots, y_m)$ (where the

[3] If m is the number of locations of the transition system, $\bar{Y}$ ranges over $\bigl(\wp(\mathbb{R}^n)\bigr)^m$ and $\bar{F} \colon \bigl(\wp(\mathbb{R}^n)\bigr)^m \to \bigl(\wp(\mathbb{R}^n)\bigr)^m$. Similarly, $\bar{Z}$ ranges over A^m and $\bar{G} \colon A^m \to A^m$.

Q_i's are first-order quantifiers $\forall$ and $\exists$, and ψ is a logical combination using $\wedge$, $\vee$, $\neg$ of polynomial inequalities in the variables $x_1, \ldots, x_n, y_1, \ldots, y_m$), it can be computed an equivalent quantifier-free formula $\psi'(x_1, \ldots, x_n)$ that depends only on the free variables. In particular, the first-order theory of polynomial inequalities over the reals is decidable.

It is therefore possible to express the (abstract) semantics of transition systems in terms of semi-algebraic sets (described by quantifier-free formulas) as follows:

Union. Given two formulas ψ and φ, the union of the respective semi-algebraic sets corresponds to the formula $\psi \vee \varphi$.

Intersection. Given two formulas ψ and φ, the intersection of the respective semi-algebraic sets corresponds to the formula $\psi \wedge \varphi$.

Update. Given the formula $\psi(\boldsymbol{x})$ and the update relation $\boldsymbol{x}' \in \rho(\boldsymbol{x})$ that associates the old values $\boldsymbol{x}$ to the new values $\boldsymbol{x}'$, expressed in terms of polynomial inequalities, the semantics of ρ on ψ is precisely captured by the quantifier-free formula equivalent to $\exists \boldsymbol{y} \,.\, \psi(\boldsymbol{y}) \wedge \big(\boldsymbol{x} \in \rho(\boldsymbol{y})\big)$.

Test for Inclusion. Given the formulas ψ and φ, the semi-algebraic set determined by ψ is included in the semi-algebraic set determined by φ if and only if $\psi \implies \varphi$ is valid. This can be decided by checking that $\exists \boldsymbol{y} \,.\, \psi \wedge \neg\varphi$ is false, where $\boldsymbol{y}$ is the tuple of all variables occurring in ψ and φ.

So, in order to generate polynomial inequalities as invariants, one could apply abstract interpretation using the operations above for defining the abstract semantics. However, this approach suffers from very serious limitations:

1. The domain of semi-algebraic sets has infinite ascending chains and termination can only be guaranteed by using a widening operator. However, as far as we know, no suitable widening operator has been defined on this domain. Moreover, the expressive power of the language of polynomial inequalities, seems to make it difficult to generalize the ideas that are at the basis of widenings on convex polyhedra [2, 3, 15]. Furthermore, the notion of *dimension* (employed in [32] for polynomial *equalities*) seems not to be useful in this context.
2. Although there have been major improvements on the original algorithm for quantifier elimination (whose complexity cannot be bounded by any finite tower of exponentials), the worst-case complexity of the current techniques like CAD [8] is doubly exponential in the number of variables. Thus it cannot be expected that the method scales up to even medium-sized systems.

As a consequence of these limitations, we turn our attention to techniques that offer a better trade-off between efficiency and precision. The fundamental ideas we follow in this paper are to restrain ourselves to basic semi-algebraic sets (i.e., finite conjunctions of polynomial inequalities) and to bound the degree of the polynomials.

4 Approximating Basic Semi-algebraic Sets by Polynomial Cones

The construction of our abstract domain is analogous to that in [9], where *pseudo-ideals* of polynomials are introduced to infer polynomial *equalities* as invariants, while still reasoning in the framework of linear algebra. Here, we extend this approach so as to handle polynomial *inequalities* as invariants.

In [9], the basic underlying definition is that of a *vector space* of polynomials:

Definition 6. (Vector space.) *A set of polynomials $V \subseteq \mathbb{R}[\boldsymbol{x}]$ is a* vector space *if* (1) $0 \in V$*; and* (2) $\lambda p + \mu q \in V$ *whenever* $p, q \in V$ *and* $\lambda, \mu \in \mathbb{R}$*. For each $Q \subseteq \mathbb{R}[\boldsymbol{x}]$, the vector space spanned by Q, denoted by $\mathcal{V}(Q)$, is the least vector space containing Q, that is,*

$$\mathcal{V}(Q) := \left\{ \sum_{i=1}^{s} \lambda_i q_i \in \mathbb{R}[\boldsymbol{x}] \;\middle|\; s \in \mathbb{N}, \forall i \in \{1, \ldots, s\} : \lambda_i \in \mathbb{R}, q_i \in Q \right\}.$$

Given a vector space V, we associate the constraint $p = 0$ to any $p \in V$. Notice that, if $p, q \in \mathbb{R}[\boldsymbol{x}]$ and $\boldsymbol{v} \in \mathbb{R}^n$ are such that $p(\boldsymbol{v}) = 0$ and $q(\boldsymbol{v}) = 0$, then $(\lambda p + \mu q)(\boldsymbol{v}) = 0$, for any $\lambda, \mu \in \mathbb{R}$. Further, for any $\boldsymbol{v} \in \mathbb{R}^n$, the zero polynomial trivially satisfies $0(\boldsymbol{v}) = 0$. Thus, the set of polynomials that evaluate to 0 on a set of states $S \subseteq \mathbb{R}^n$, that is $\{ p \in \mathbb{R}[\boldsymbol{x}] \mid \forall \boldsymbol{v} \in S : p(\boldsymbol{v}) = 0 \}$, has the structure of a vector space. Unfortunately, this vector space has infinite dimension.[4] In order to work with objects of finite dimension, it is necessary to approximate by bounding the degrees of the polynomials.

Moreover, when considering polynomials as elements of a vector space, the algebraic relationships between terms such as x_1, x_2 and $x_1 x_2$ are lost. For instance, consider the vector space $\mathcal{V}\bigl(\{x_1, x_2 - x_1 x_2\}\bigr)$, generated by the polynomial equalities $x_1 = 0$ and $x_2 = x_1 x_2$. Then, even though the polynomial equality $x_2 = 0$ is semantically entailed by the previous ones, $x_2 \notin \mathcal{V}\bigl(\{x_1, x_2 - x_1 x_2\}\bigr)$. The reason is that the vector space generated by x_1 and $x_2 - x_1 x_2$ only includes the *linear* combinations of its generators, whereas in the case above x_2 can only be obtained by a *nonlinear* combination of the generators, namely $x_2 = x_2 \cdot (x_1) + 1 \cdot (x_2 - x_1 x_2)$. This problem can be solved by adding the polynomial $x_1 x_2$ to the set of generators, so that the polynomial $x_2 \in \mathcal{V}\bigl(\{x_1, x_1 x_2, x_2 - x_1 x_2\}\bigr)$ can be obtained by the linear combination $0 \cdot (x_1) + 1 \cdot (x_1 x_2) + 1 \cdot (x_2 - x_1 x_2)$.

In general, in order to reduce the loss of precision due to the linearization of the abstraction, additional polynomials are added taking into account that, when $p \in \mathbb{R}[\boldsymbol{x}]$ and $\boldsymbol{v} \in \mathbb{R}^n$ are such that $p(\boldsymbol{v}) = 0$, we have $(pq)(\boldsymbol{v}) = 0$ for each $q \in \mathbb{R}[\boldsymbol{x}]$. Therefore, pseudo-ideals are defined as follows:

Definition 7. (Pseudo-ideal.) *Given a degree bound $d \in \mathbb{N}$, a* pseudo-ideal *$P \subseteq \mathbb{R}_d[\boldsymbol{x}]$ of degree d is a vector space with the property that $pq \in P$ whenever $p \in P$, $q \in \mathbb{R}[\boldsymbol{x}]$ and $\deg(pq) \leq d$. For each $Q \subseteq \mathbb{R}_d[\boldsymbol{x}]$, the pseudo-ideal of degree d spanned by Q, denoted by $\mathcal{P}_d(Q)$, is the least pseudo-ideal of degree d containing Q.*

[4] Actually it has the algebraic structure of an *ideal* of polynomials [17].

Pseudo-ideals are "complete" in the sense that they are closed under the operations of addition, product by scalars and bounded product by polynomials. For instance, $x_1x_2 \in \mathcal{P}_2(x_1, x_2)$. Pseudo-ideals are the elements of the abstract domain used in [9].

In order to extend this methodology to the generation of invariant polynomial inequalities, a first, necessary step is the identification, in the basic semi-algebraic context, of an adequate algebraic structure playing the same role of vector spaces for polynomial equalities. It turns out that *polynomial cones* are the right notion:

Definition 8. (Polynomial cone.) *A set of polynomials $C \subseteq \mathbb{R}[\boldsymbol{x}]$ is a* polynomial cone *if* (1) $1 \in C$*; and* (2) $\lambda p + \mu q \in C$ *whenever* $p, q \in C$ *and* $\lambda, \mu \in \mathbb{R}_+$. *For each $Q \subseteq \mathbb{R}[\boldsymbol{x}]$, the* polynomial cone generated by Q*, denoted by $\mathcal{C}(Q) \subseteq \mathbb{R}[\boldsymbol{x}]$, is the least polynomial cone containing Q, that is,*

$$\mathcal{C}(Q) := \left\{ \lambda + \sum_{i=1}^{s} \lambda_i q_i \in \mathbb{R}[\boldsymbol{x}] \;\middle|\; \lambda \in \mathbb{R}_+, s \in \mathbb{N}, \forall i \in \{1, \dots, s\} : \lambda_i \in \mathbb{R}_+, q_i \in Q \right\}.$$

Mimicking the reasoning done before for vector spaces, we associate the constraint $p \geq 0$ to any polynomial p in the polynomial cone C. Consider the basic semi-algebraic set defined by the constraint system

$$\psi = \{f_1 = 0, \dots, f_h = 0, g_1 \geq 0, \dots, g_k \geq 0\} \tag{2}$$

where, for each $i = 1, \dots, h$ and $j = 1, \dots, k$, we have $f_i, g_j \in \mathbb{R}[\boldsymbol{x}]$. Then, the set of polynomial inequalities that are consequences of ψ define a polynomial cone. Indeed, $\psi \implies (1 \geq 0)$ trivially; and, if $\psi \implies (p \geq 0)$ and $\psi \implies (q \geq 0)$, clearly $\psi \implies (\lambda p + \mu q \geq 0)$ for each $\lambda, \mu \in \mathbb{R}_+$. As was the case for the vector space of polynomials, this set of polynomials has infinite dimension. In order to deal with objects of finite dimension and to mitigate the precision loss due to linearization, we again fix an upper bound for the degrees of the polynomials and then complete with respect to bounded product by polynomials. The analog of pseudo-ideals in the basic semi-algebraic setting are *complete polynomial cones*:

Definition 9. (Complete polynomial cone.) *Given a degree bound $d \in \mathbb{N}$, a* complete polynomial cone *of degree d is a polynomial cone $C \subseteq \mathbb{R}_d[\boldsymbol{x}]$ satisfying:*

(1) *$pq \in C$ whenever $p, q \in C$ and* $\deg(pq) \leq d$*;*
(2) *$pq \in C$ whenever $p, -p \in C$, $q \in \mathbb{R}[\boldsymbol{x}]$ and* $\deg(pq) \leq d$.

For each $Q \subseteq \mathbb{R}_d[\boldsymbol{x}]$, the complete polynomial cone of degree d generated by Q*, denoted by $\mathcal{C}_d(Q) \subseteq \mathbb{R}_d[\boldsymbol{x}]$, is the least complete polynomial cone of degree d containing Q.*

Let ψ be a constraint system defining a basic semi-algebraic set. Then, once the degree bound $d \in \mathbb{N}$ is fixed, ψ is abstracted by computing the complete polynomial cone $C := \mathcal{C}_d\bigl(\mathrm{poly}(\psi) \cap \mathbb{R}_d[\boldsymbol{x}]\bigr)$. The approximation forgets those polynomials occurring in ψ having a degree greater than d. Also note that the precision of the approximation depends on the specific constraint system.[5] Consider the constraint system ψ as defined in (2). Let $\mathcal{M}(g_1, \ldots, g_k)$ be the *multiplicative monoid* generated by the g_j's, i.e., the set of finite products of g_j's including 1 (the empty product). From the definition, it can be seen that the polynomials $p \in C = \mathcal{C}_d\bigl(\mathrm{poly}(\psi) \cap \mathbb{R}_d[\boldsymbol{x}]\bigr)$ can all be obtained as[6]

$$p = \sum_{i=1}^{h} r_i f_i + \sum_j \lambda_j q_j, \tag{3}$$

where, for each $i = 1, \ldots, h$, $r_i \in \mathbb{R}[\boldsymbol{x}]$ is such that $\deg(r_i f_i) \leq d$, and, for each j, $\lambda_j \in \mathbb{R}_+$ and $q_j \in \mathcal{M}(g_1, \ldots, g_k) \cap \mathbb{R}_d[\boldsymbol{x}]$.

The abstraction is sound, since for each $p \in C$ and each $\boldsymbol{v} \in \mathbb{R}^n$ such that $\psi(\boldsymbol{v})$ holds, we have

$$p(\boldsymbol{v}) = \sum_{i=1}^{h} r_i(\boldsymbol{v}) f_i(\boldsymbol{v}) + \sum_j \lambda_j q_j(\boldsymbol{v}) = \sum_j \lambda_j q_j(\boldsymbol{v}) \geq 0.$$

The abstraction is also complete for the linear case. In fact, consider any finite set of linear constraints $\varphi = \{p_1 \geq 0, \ldots, p_m \geq 0\}$, which we assume to be satisfiable. Then, the corresponding complete polynomial cone of degree 1 is $L = \mathcal{C}_1\bigl(\mathrm{poly}(\varphi)\bigr)$, whose elements are linear consequences of φ. On the other hand, if $p \in \mathbb{R}_1[\boldsymbol{x}]$ is a linear polynomial such that $\varphi \implies (p \geq 0)$, then by Farkas' lemma there exists $\boldsymbol{\mu} = (\mu_0, \ldots, \mu_m) \in \mathbb{R}_+^{m+1}$ such that $p = \mu_0 + \sum_{i=1}^{m} \mu_i p_i$; in other words, $p \in L$. In the general nonlinear setting, the abstraction constituted by complete polynomial cones is not complete. Notice however that the set of *all* invariant polynomial inequalities (i.e., invariants of the form $p \geq 0$ with $p \in \mathbb{R}[\boldsymbol{x}]$) is not computable in basic semi-algebraic transition systems [29]. Worse, the set of all invariant *linear equalities* (i.e., invariants of the form $p = 0$ with $p \in \mathbb{R}_1[\boldsymbol{x}]$) is not computable in transition systems even if restricted to linear equality guards. Therefore, a complete abstraction (such as the one sketched in Section 3) still would not allow to obtain a complete method for generating invariant polynomial inequalities.

4.1 Representation

Just as linear algebra is used in [9] to manipulate vector spaces representing pseudo-ideals of degree d, in this paper we exploit the theory of convex polyhedra to handle (finitely generated) polynomial cones representing complete polynomial cones of degree d.

The linearization in the abstraction process implies that all terms are considered as different variables. For instance, in Example 1, the terms a, b, c, c^2 are

[5] For instance, the equivalent constraint systems $\{x = 0\}$ and $\{x^2 = 0\}$ are abstracted to different complete polynomial cones of degree 1.

[6] It is further necessary that the following non-degeneracy condition is satisfied: there exists $\boldsymbol{v} \in \mathbb{R}^n$ such that $\psi(\boldsymbol{v})$ holds and $\prod_{j=1}^{k} g_j(\boldsymbol{v}) > 0$.

```
Require: A finite set of polynomial equalities $\varphi = \{f_1 = 0, \dots, f_h = 0\}$ and a finite
  set of polynomial inequalities $\psi = \{g_1 \geq 0, \dots, g_k \geq 0\}$.
Ensure: $\varphi' = \{f'_1 = 0, \dots, f'_{h'} = 0\}$ and $\psi' = \{g'_1 \geq 0, \dots, g'_{k'} \geq 0\}$ are finite sets
  of polynomial equalities and inequalities, respectively, such that $\mathcal{C}\bigl(\mathrm{poly}(\varphi' \cup \psi')\bigr) =$
  $\mathcal{C}_d\bigl(\mathrm{poly}(\varphi \cup \psi) \cap \mathbb{R}_d[\boldsymbol{x}]\bigr)$.

  $\varphi' := \psi' := \emptyset$
  for all $(f = 0) \in \varphi$ do
    if $\deg(f) \leq d$ then
      for all $\boldsymbol{x}^{\boldsymbol{\alpha}}$ such that $\deg(\boldsymbol{x}^{\boldsymbol{\alpha}}) \leq d - \deg(f)$ do
        $\varphi' := \varphi' \cup \{\boldsymbol{x}^{\boldsymbol{\alpha}} f = 0\}$
  for all finite product $g^*$ of $g$'s such that $(g \geq 0) \in \psi$ do
    if $\deg(g^*) \leq d$ then
      $\psi' := \psi' \cup \{g^* \geq 0\}$
```

Fig. 2. Algorithm $\mathrm{complete}_d$

all regarded as different and potentially independent variables, so that the initial condition $\mathcal{I}(\ell_0) = \{a \geq 0, b = 0, c \geq 1, c^2 \geq a + 1\}$ is interpreted as defining a convex polyhedron in an ambient space of dimension at least 4. In general, given a transition system on an n-tuple $\boldsymbol{x}$ of variables and a degree bound d, the introduction of the auxiliary variables, standing for all the nonlinear terms of degree at most d, yields an m-tuple $\boldsymbol{y}$ of variables, where each y_i corresponds to one of the $m = \binom{n+d}{d} - 1$ different terms $\boldsymbol{x}^{\boldsymbol{\alpha}} \in \mathbb{R}_d[\boldsymbol{x}] \setminus \{1\}$. In the following, we will denote each y_i by writing the corresponding term. Computation in the abstract domain of cones of degree d is feasible provided d is small, e.g., 2 or 3. We have implemented the techniques presented in this paper for the case $d = 2$; in Section 5 we show the results obtained with this implementation.

It remains to be seen how the linearized constraint system can be completed according to Definition 9. Algorithm $\mathrm{complete}_d$, which is based on Equation (3), is given in Figure 2.

Example 2. Consider Example 1. The completion of degree 2 of the polynomial cone corresponding to the initial condition $\mathcal{I}(\ell_0)$ yields the system of constraints

$$\begin{aligned}
\mathcal{C}_2\bigl(\mathcal{I}(\ell_0)\bigr) &= \mathcal{C}_2\bigl(\{b = 0\} \cup \{a \geq 0, c \geq 1, c^2 \geq a + 1\}\bigr) \\
&= \{b = 0, ab = 0, b^2 = 0, bc = 0\} \\
&\qquad \cup \{a \geq 0, c \geq 1, c^2 \geq a + 1, a^2 \geq 0, c^2 \geq 1, ac \geq 0\}.
\end{aligned}$$

4.2 Abstract Semantics

In this section we review the operations required in order to perform abstract interpretation of transition systems using polynomial cones as abstract values.

Union. Given two (finitely generated) polynomial cones C_1 and C_2 representing the polynomial constraint systems ψ_1 and ψ_2, respectively, we would like to approximate the union of the corresponding basic semi-algebraic sets using another basic semi-algebraic set. By duality, this amounts to computing the intersection cone $C_1 \cap C_2$: for each $p \in C_1 \cap C_2$ and $\boldsymbol{v} \in \mathbb{R}^n$ such that $\psi_1(\boldsymbol{v}) \vee \psi_2(\boldsymbol{v})$, either $\psi_1(\boldsymbol{v})$, so that $p(\boldsymbol{v}) \geq 0$ as $p \in C_1$; or $\psi_2(\boldsymbol{v})$, so that $p(\boldsymbol{v}) \geq 0$ as $p \in C_2$. Thus, the approximation is sound. At the implementation level, using the representation of polynomial cones as convex polyhedra, this intersection of cones corresponds to the convex polyhedral hull operation.

Intersection. Given two (finitely generated) polynomial cones $C_1 = \mathcal{C}(Q_1)$ and $C_2 = \mathcal{C}(Q_2)$, we would like to compute the intersection of the respective basic semi-algebraic sets. Then a sound approximation is to compute the cone spanned by the union of the generators, $\mathcal{C}(Q_1 \cup Q_2)$. In order to reduce the loss of precision due to linearization, we complete this cone up to the degree bound d. Thus, the polynomial cone corresponding to the intersection is $\mathrm{complete}_d(Q_1 \cup Q_2)$.

Update. Each (basic semi-algebraic) update map $\rho\colon \mathbb{R}^n \to \wp(\mathbb{R}^n)$, working on the original n-tuple of variables $\boldsymbol{x}$, is approximated by an affine update map $\rho'\colon \mathbb{R}^m \to \wp(\mathbb{R}^m)$, where $m = \binom{n+d}{d} - 1$, working on the extended m-tuple $\boldsymbol{y}$ of terms. The new update map ρ' is obtained by composing a sequence of simpler affine maps, each one approximating the effect of ρ on a single term. For the sake of notation, if variable y_i corresponds to term $\boldsymbol{x}^{\boldsymbol{\alpha}}$ and $p \in \mathbb{R}_d[\boldsymbol{x}]$, let $\boldsymbol{x}^{\boldsymbol{\alpha}} \mapsto p$ denote the update map such that, for each $\boldsymbol{w} \in \mathbb{R}^m$,

$$(\boldsymbol{x}^{\boldsymbol{\alpha}} \mapsto p)(\boldsymbol{w}) := \bigl(w_1, \dots, w_{i-1}, p(\boldsymbol{w}), w_{i+1}, \dots, w_m\bigr) \in \mathbb{R}^m. \tag{4}$$

Note that the (possibly nonlinear) polynomial $p \in \mathbb{R}_d[\boldsymbol{x}]$ on the original tuple of variables is interpreted as a linear polynomial $p \in \mathbb{R}_1[\boldsymbol{y}]$ on the extended ambient space, so that Equation (4) indeed defines an affine map.

By hypothesis, ρ is defined by $k \leq n$ polynomials $p_1, \dots, p_k \in \mathbb{R}[\boldsymbol{x}]$ and distinct indices $i_1, \dots, i_k \in \{1, \dots, n\}$ such that, for each $\boldsymbol{v} \in \mathbb{R}^n$,

$$\rho(\boldsymbol{v}) = \bigl\{\, (v'_1, \dots, v'_n) \in \mathbb{R}^n \bigm| v'_{i_1} = p_1(\boldsymbol{v}), \dots, v'_{i_k} = p_k(\boldsymbol{v}) \,\bigr\}.$$

Then, for each term $\boldsymbol{x}^{\boldsymbol{\alpha}} \in \mathbb{R}_d[\boldsymbol{x}] \setminus \{1\}$, we distinguish the following cases:

- Suppose there exists $j \in \{1, \dots, n\}$ such that $\alpha_j > 0$ and $j \notin \{i_1, \dots, i_k\}$. This means that ρ nondeterministically updates at least one of the relevant factors of the term $\boldsymbol{x}^{\boldsymbol{\alpha}}$. Thus, we conservatively approximate the overall effect of ρ on $\boldsymbol{x}^{\boldsymbol{\alpha}}$ as if it was a nondeterministic assignment.
- Suppose now that, for each $j = 1, \dots, n$, if $\alpha_j > 0$ then $j \in \{i_1, \dots, i_k\}$, i.e., all the relevant factors of $\boldsymbol{x}^{\boldsymbol{\alpha}}$ are deterministically updated by ρ. Then:
  - if the polynomial $p_{\boldsymbol{\alpha}} := \prod\bigl\{\, p_h^{\alpha_j}(\boldsymbol{x}) \bigm| j \in \{1, \dots, n\}, \alpha_j > 0, j = i_h \,\bigr\}$ is such that $p_{\boldsymbol{\alpha}} \in \mathbb{R}_d[\boldsymbol{x}]$, we apply the affine map $\boldsymbol{x}^{\boldsymbol{\alpha}} \mapsto p_{\boldsymbol{\alpha}}$;
  - otherwise, since we cannot represent the effect of ρ on $\boldsymbol{x}^{\boldsymbol{\alpha}}$, we (again) conservatively overapproximate it as a nondeterministic assignment.

Since ρ updates all terms simultaneously, these maps are ordered topologically according to the dependencies of terms (possibly adding temporary copies of some term variables, which are eliminated at the end).

Example 3. Consider the transitions of Example 1. For the transition τ_0 we have $\rho_0 \equiv (a', b', c') = (a, b, c/2)$. This update is encoded by the sequence of affine maps $ac \mapsto ac/2$, $bc \mapsto bc/2$, $c^2 \mapsto c^2/4$ and $c \mapsto c/2$, leading to ρ_0' defined as

$$\bigl(a', b', c', ab', ac', bc', (a^2)', (b^2)', (c^2)'\bigr) = (a, b, c/2, ab, ac/2, bc/2, a^2, b^2, c^2/4).$$

Test for inclusion. The test for inclusion can be conservatively overapproximated by means of the test for inclusion for convex polyhedra.

Widening. Any widening for convex polyhedra, e.g., the standard widening [15] or the more sophisticated widenings proposed in [2, 3], will serve the purpose of guaranteeing termination, with different trade-offs between efficiency and precision.

Example 4. For the transitions of Example 1, using the abstract semantics shown above, we obtain the invariant

$$\begin{aligned} \mathrm{reach}_{\ell_0}(S) \implies \bigl\{&(b+c)^2 \geq a+1, a \geq b^2, b \geq 0, c \geq 1, \\ &a^2 \geq 0, ab \geq 0, ac \geq 0, b^2 \geq bc, bc \geq b, (c-1)^2 \geq 0\bigr\}. \end{aligned}$$

Notice that all the constraints appearing on the second line are in fact redundant. Some of these, such as $(c-1)^2 \geq 0$ and $a^2 \geq 0$, are trivially redundant in themselves. Other ones are made redundant by the constraints appearing on the first line (for instance, $ab \geq 0$ is implied by $a \geq b^2$ and $b \geq 0$). This phenomenon is due to the interaction of the completion procedure, which adds redundant constraints to polynomial cones, with the underlying linear inequalities inference rules, which are treating different terms as independent variables and, as a consequence, are only able to detect and remove some of the redundancies.

The two constraints $(b+c)^2 \geq a+1$ and $a \geq b^2$ in the invariant above are essential in a formal proof of the (partial) correctness of the program in Figure 1. Note that the computed invariant *assumes* that the integer division $c := c/2$ is correctly modeled by rational division. Such an assumption can be validated by other analyses, e.g., by using a domain of numerical powers [26], which could infer that c evaluates to a power of 2 at location ℓ_0. Since on termination $c = 1$ holds, the conjunction of these constraints implies $(b+1)^2 > a \geq b^2$.

5 Experimental Evaluation

The approach described in this paper has been implemented in a prototype analyzer that infers polynomial inequalities of degree not greater than $d = 2$. The prototype, which is based on the *Parma Polyhedra Library* (PPL) [4], first performs a rather standard *linear* relations analysis, then assumes the linear invariants so obtained for the analysis of (possibly) nonlinear invariants described in the previous sections. We have observed that this preliminary linear analysis improves the results in a significant way. In fact:

1. it ensures that we never obtain less information than is achievable with the linear analysis alone;

2. the availability of “trusted” linear invariants increases the precision of the nonlinear analysis considerably;
3. the time spent in the linear analysis phase is usually recovered in the quadratic analysis phase.

The prototype uses the sophisticate widening operator proposed in [2] enhanced with variations of the “widening up to” technique described in [21] and with the “widening with tokens” technique (a form of delayed widening application) described in [3].

Considering that, with the chosen degree bound $d = 2$, we are working on an ambient space that has a dimension which is quadratic in the number of variables of the transition system being analyzed, and considering that polyhedra operations have exponential worst-case complexity, some care has to be taken in order to analyze systems of realistic complexity. In our prototype, we exploit the capability of the PPL concerning the support of time-bounded computations. All polyhedra operations are subject to a timeout (5 seconds of CPU time in the experimentation we are about to report); when a timeout expires, the PPL abandons (without leaking memory) the current computation and gives control back to the analyzer. This situation is handled by the analyzer by using a less precise operation (such as replacing the precise convex polyhedral hull of two polyhedra P_1 and P_2 by the polyhedron obtained by removing, from a system of constraints defining P_1, all constraints that are not satisfied by P_2) or by simplifying the involved polyhedra resorting to a domain of bounded differences.[7] With this technique we are able to obtain results that are generally quite precise in reasonable time (note that the prototype was not coded with speed in mind).

We have run the prototype analyzer on a benchmark suite constituted by all the programs from the FAST suite [5] (`http://www.lsv.ens-cachan.fr/fast/`), programs taken from the StInG suite [34] (`http://www.stanford.edu/~srirams/Software/sting.html`), all square root algorithms in [18], programs from [10, 22, 33, 39], and a program, `array`, written by the authors. From the StInG suite we have only omitted those programs with nondeterministic assignments where the *rvalue* is bounded by linear expressions (like $0 \geq x' \geq x + y$), which do not fall into the programming model used here.

A summary of the experimental results is presented in Table 1. Besides the program name, its origin and the number of variables, locations and transitions (columns from 1 to 5, respectively), the table indicates: (1) the CPU time, in seconds, taken to compute our *linear* invariants (column 6) and how they compare with the ones computed by StInG (column 7: ‘+’ means ours are better, ‘−’ means ours are worse, ‘=’ means they are equal, ‘≠’ means they are not comparable); and (2) the time taken to generate *quadratic* invariants (column 8) and whether these invariants improve upon (that is, are not implied by) the linear ones, *taking into account both our linear invariants as well as those generated by StInG* (column 9: ‘✓’ means we improve the precision). The measurements

[7] The reason our current prototype does not resort to the more precise domain of octagons is contingent and only due to the fact that the implementation of octagons in the PPL is not yet ready for production use.

Table 1. A summary of the experimental results

					Linear analysis		Quadratic analysis	
Program name	Origin	n	$\|\mathcal{L}\|$	$\|\mathcal{T}\|$	CPU time	vs StInG	CPU time	Improves
`array`		4	5	6	0.2	+ ≠≠≠ +	79.8	✓
`bakery`	[39]	2	9	24	18.6	= ··· =	0.2	✓
`barber`	FAST	8	1	12	18.7	−	2.7	✓
`berkeley`	FAST	4	1	3	0.0	+	0.1	✓
`cars`	StInG	7	1	2	18.5	≠	45.9	✓
`centralserver`	FAST	12	1	8	5.4	+	193.4	
`consistency`	FAST	11	1	7	2.5	=	10.0	
`consprodjava`	FAST	16	1	14	325.6	+	601.9	
`consprodjavaN`	FAST	16	1	14	308.0	+	611.6	
`cousot05vmcai`	[10]	4	1	1	0.0	=	0.1	✓
`csm`	FAST	14	1	13	29.3	=	219.5	
`dekker`	FAST	22	1	22	458.4	=	1218.1	
`dragon`	FAST	5	1	12	0.5	−	1.4	✓
`efm`	FAST	6	1	5	0.1	=	0.3	
`rfm05hscc`	[33]	4	1	2	0.1	≠	38.5	
`firefly`	FAST	4	1	8	0.1	=	0.2	✓
`fms`	FAST	22	1	20	893.2	=	2795.0	
`freire`	[19]	3	1	1	0.0	−	6.4	
`futurbus`	FAST	9	1	9	2.8	+	23.2	✓
`heap`	StInG	5	1	4	0.1	≠	10.9	
`illinois`	FAST	4	1	9	0.1	=	0.3	✓
`kanban`	FAST	16	1	16	60.5	=	340.4	
`lamport`	FAST	11	1	9	3.1	+	13.4	
`lifo`	StInG	7	1	10	1.4	+	14.8	✓
`lift`	FAST	4	1	5	0.1	=	22.1	
`mesi`	FAST	4	1	4	0.0	=	0.1	✓
`moesi`	FAST	5	1	4	0.1	−	0.3	✓
`multipoll`	FAST	18	1	17	116.3	=	476.8	
`peterson`	FAST	14	1	12	17.6	+	88.5	
`producer-consumer`	FAST	5	1	3	0.1	=	15.5	
`readwrit`	FAST	13	1	9	7.7	=	2147.3	
`rtp`	FAST	9	1	12	2.6	=	8.9	
`see-saw`	StInG	2	1	4	0.0	−	5.3	
`sqroot1`	[18]	2	1	1	0.0	=	0.0	✓
`sqroot2`	[18]	3	1	8	0.0	+	15.6	✓
`sqroot3`	[18]	3	2	6	0.0	==	10.3	✓
`sqroot4`	[18]	4	2	6	10.3	=+	8.2	✓
`sqroot5`	[9]	4	1	1	0.0	+	6.1	✓
`sqroot6`	[22]	5	1	2	0.0	=	15.5	✓
`swim-pool`	StInG	9	1	6	1.5	+	46.5	
`synapse`	FAST	3	1	3	0.0	+	0.0	✓
`ticket2i`	FAST	6	1	6	0.3	+	5.8	
`ticket3i`	FAST	8	1	9	9.5	+	82.6	
`train-beacon`	StInG	3	4	12	0.1	= −− =	20.5	
`train-one-loc`	StInG	3	1	6	0.0	−	0.4	
`ttp`	FAST	9	4	17	9.3	++++	126.9	

were performed on a PC with an Intel® Xeon$^{\mathrm{TM}}$ CPU clocked at 1.80 GHz, equipped with 1 GB of RAM and running GNU/Linux. Notice that for about 80% of the locations, our linear invariants are at least as strong as the ones produced by StInG, and that, in fact, for one third ours are stronger. Most importantly, for about half of the programs, the obtained quadratic invariants improve the precision of the linear analysis.

6 Conclusion

We have presented a technique for generating invariant polynomial inequalities of bounded degree. The technique, which is based on the abstract interpretation framework, consists in overapproximating basic semi-algebraic sets by means of convex polyhedra, and can thus take advantage of all the work done in that field (e.g., refined widening operators, devices able to throttle the complexity of the analysis such as restricted classes of polyhedra, ways of partitioning the vector space and so forth). The application of our prototype implementation to a number of benchmark programs shows that the method can produce non-trivial and useful quadratic invariant inequalities in reasonable time, thus proving the feasibility of the automatic inference of nonlinear invariant inequalities (something that was previously unclear).

For future work, we want to generalize our definition of basic semi-algebraic transition system so as to capture a form of nondeterministic assignments where the *rvalue* is bounded by means of polynomial inequalities, rather than being completely unknown. We would also like to increase the precision of the approach by incorporating, in the $\mathrm{complete}_d$ algorithm, other forms of inference, such as *relational arithmetic* [1, 37]. This technique allows to infer constraints on the qualitative relationship of an expression to its arguments and can be expressed by a number of axiom schemata such as

$$(x > 0 \land y > 0) \implies \bigl(x \bowtie 1 \implies xy \bowtie y\bigr),$$

which is valid for each $\bowtie \in \{=, \neq, \leq, <, \geq, >\}$. Finally, there is much room for improving the prototype implementation. To start with, we believe its performance can be greatly enhanced (there are a number of well-known techniques that we are not currently using); this may even bring us to the successful inference of cubic invariants for simple programs. The simplification of the analysis results is another natural candidate for this line of work.

Acknowledgments. The authors thank D. Kapur for his help and comments.

References

1. R. Bagnara. *Data-Flow Analysis for Constraint Logic-Based Languages.* PhD thesis, Dipartimento di Informatica, Università di Pisa, Pisa, Italy, March 1997. Printed as Report TD-1/97.

2. R. Bagnara, P. M. Hill, E. Ricci, and E. Zaffanella. Precise widening operators for convex polyhedra. In R. Cousot, editor, *Static Analysis: Proceedings of the 10th International Symposium*, volume 2694 of *Lecture Notes in Computer Science*, pages 337–354, San Diego, California, USA, 2003. Springer-Verlag, Berlin.
3. R. Bagnara, P. M. Hill, E. Ricci, and E. Zaffanella. Precise widening operators for convex polyhedra. *Science of Computer Programming*, 2005. To appear.
4. R. Bagnara, E. Ricci, E. Zaffanella, and P. M. Hill. Possibly not closed convex polyhedra and the Parma Polyhedra Library. In M. V. Hermenegildo and G. Puebla, editors, *Static Analysis: Proceedings of the 9th International Symposium*, volume 2477 of *Lecture Notes in Computer Science*, pages 213–229, Madrid, Spain, 2002. Springer-Verlag, Berlin.
5. S. Bardin, A. Finkel, J. Leroux, and L. Petrucci. FAST: Fast Acceleration of Symbolic Transition systems. In W. A. Hunt, Jr. and F. Somenzi, editors, *Computer Aided Verification: Proceedings of the 15th International Conference*, volume 2725 of *Lecture Notes in Computer Science*, pages 118–121, Boulder, CO, USA, 2003. Springer-Verlag, Berlin.
6. S. Bensalem, M. Bozga, J.-C. Fernandez, L. Ghirvu, and Y. Lakhnech. A transformational approach for generating non-linear invariants. In J. Palsberg, editor, *Static Analysis: 7th International Symposium, SAS 2000*, volume 1824 of *Lecture Notes in Computer Science*, pages 58–74, Santa Barbara, CA, USA, 2000. Springer-Verlag, Berlin.
7. 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'03)*, pages 196–207, San Diego, California, USA, 2003. ACM Press.
8. G.E. Collins. Quantifier Elimination for Real Closed Fields by Cylindrical Algebraic Decomposition. In H. Brakhage, editor, *Automata Theory and Formal Languages*, volume 33 of *LNCS*, pages 134–183. Springer-Verlag, 1975.
9. M. Colón. Constraint-based linear-relations analysis. In Giacobazzi [20], pages 296–311.
10. P. Cousot. Proving program invariance and termination by parametric abstraction, lagrangian relaxation and semidefinite programming. In Cousot [16], pages 1–24.
11. P. Cousot and R. Cousot. Static determination of dynamic properties of programs. In B. Robinet, editor, *Proceedings of the Second International Symposium on Programming*, pages 106–130. Dunod, Paris, France, 1976.
12. P. Cousot and R. Cousot. Abstract interpretation: A unified lattice model for static analysis of programs by construction or approximation of fixpoints. In *Proceedings of the Fourth Annual ACM Symposium on Principles of Programming Languages*, pages 238–252, New York, 1977. ACM Press.
13. P. Cousot and R. Cousot. Systematic design of program analysis frameworks. In *Proceedings of the Sixth Annual ACM Symposium on Principles of Programming Languages*, pages 269–282, New York, 1979. ACM Press.
14. P. Cousot, R. Cousot, J. Feret, L. Mauborgne, A. Miné, D. Monniaux, and X. Rival. The ASTRÉE analyzer. In M. Sagiv, editor, *Programming Languages and Systems, Proceedings of the 14th European Symposium on Programming*, volume 3444 of *Lecture Notes in Computer Science*, pages 21–30, Edinburgh, UK, 2005. Springer-Verlag, Berlin.
15. P. Cousot and N. Halbwachs. Automatic discovery of linear restraints among variables of a program. In *Conference Record of the Fifth Annual ACM Symposium*

on Principles of Programming Languages, pages 84–96, Tucson, Arizona, 1978. ACM Press.

16. R. Cousot, editor. *Verification, Model Checking and Abstract Interpretation: Proceedings of the 6th International Conference (VMCAI 2005)*, volume 3385 of *Lecture Notes in Computer Science*, Paris, France, 2005. Springer-Verlag, Berlin.
17. D. A. Cox, J. B. Little, and D. O'Shea. *Ideals, Varieties and Algorithms.* Springer-Verlag, Berlin, second edition, 1996.
18. E. W. Dijkstra. *A Discipline of Programming.* Prentice-Hall, 1976.
19. P. Freire. SQRT. Retrieved April 10, 2005, from `http://www.pedrofreire.com/sqrt`, 2002.
20. R. Giacobazzi, editor. *Static Analysis: Proceedings of the 11th International Symposium*, volume 3148 of *Lecture Notes in Computer Science*, Verona, Italy, 2004. Springer-Verlag, Berlin.
21. N. Halbwachs. Delay analysis in synchronous programs. In C. Courcoubetis, editor, *Computer Aided Verification: Proceedings of the 5th International Conference*, volume 697 of *Lecture Notes in Computer Science*, pages 333–346, Elounda, Greece, 1993. Springer-Verlag, Berlin.
22. P. Hsieh. How to calculate square roots. Retrieved April 10, 2005, from `http://www.azillionmonkeys.com/qed/sqroot.html`, 2004.
23. N. D. Jones and X. Leroy, editors. *Proceedings of the 31st ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages (POPL 2004)*, Venice, Italy, 2004. ACM.
24. D. Kapur. Automatically generating loop invariants using quantifier elimination. In *Proceedings of the 10th International Conference on Applications of Computer Algebra (ACA-2004)*, Beaumont, Texas, 2004. Also published as Technical Report TR-CS-2003-58, Department of Computer Science, University of New Mexico, Albuquerque, USA, 2003.
25. M. Karr. Affine relationships among variables of a program. *Acta Informatica*, 6:133–151, 1976.
26. I. Mastroeni. Numerical power analysis. In O. Danvy and A. Filinski, editors, *Proceedings of the 2nd Symposium on Programs as Data Objects (PADO 2001)*, volume 2053 of *Lecture Notes in Computer Science*, pages 117–137, Aarhus, Denmark, 2001. Springer-Verlag, Berlin.
27. A. Miné. The octagon abstract domain. In *Proceedings of the Eighth Working Conference on Reverse Engineering (WCRE'01)*, pages 310–319, Stuttgart, Germany, 2001. IEEE Computer Society Press.
28. M. Müller-Olm and H. Seidl. Computing polynomial program invariants. *Information Processing Letters*, 91(5):233–244, 2004.
29. M. Müller-Olm and H. Seidl. A note on Karr's algorithm. In J. Diaz, J. Karhumäki, and A. Lepistö et al., editors, *Automata, Languages and Programming: Proceedings of the 31st International Colloquium (ICALP 2004)*, volume 3142 of *Lecture Notes in Computer Science*, pages 1016–1028, Turku, Finland, 2004. Springer-Verlag, Berlin.
30. M. Müller-Olm and H. Seidl. Precise interprocedural analysis through linear algebra. In Jones and Leroy [23].
31. E. Rodríguez-Carbonell and D. Kapur. An abstract interpretation approach for automatic generation of polynomial invariants. In Giacobazzi [20], pages 280–295.
32. E. Rodríguez-Carbonell and D. Kapur. Automatic generation of polynomial loop invariants: Algebraic foundations. In J. Gutierrez, editor, *Proceedings of the International Symposium on Symbolic and Algebraic Computation (ISSAC 2004)*, pages 266–273, Santander, Spain, 2004. ACM Press.

33. M. Roozbehani, E. Feron, and A. Megrestki. Modeling, optimization and computation for software verification. In M. Morari and L. Thiele, editors, *Hybrid Systems: Computation and Control: Proceedings of the 8th International Workshop (HSCC 2005)*, pages 606–622, Zürich, Switzerland, 2005.
34. S. Sankaranarayanan, H. Sipma, and Z. Manna. Constraint-based linear-relations analysis. In Giacobazzi [20], pages 53–68.
35. S. Sankaranarayanan, H. Sipma, and Z. Manna. Scalable analysis of linear systems using mathematical programming. In Cousot [16], pages 25–41.
36. S. Sankaranarayanan, H. B. Sipma, and Z. Manna. Non-linear loop invariant generation using Gröbner bases. In Jones and Leroy [23], pages 318–329.
37. R. Simmons. Commonsense arithmetic reasoning. In T. Kehler and S. Rosenschein, editors, *Proceedings of the Fifth National Conference on Artificial Intelligence (AAAI-86)*, volume 1, pages 118–124, Philadelphia, PA, 1986. AAAI Press.
38. A. Tarski. A Decision Method for Elementary Algebra and Geometry. University of California Press, 1951.
39. A. Tiwari, H. Rueß, H. Saïdi, and N. Shankar. A technique for invariant generation. In T. Margaria and W. Yi, editors, *Tools and Algorithms for Construction and Analysis of Systems, 7th International Conference, TACAS 2001*, volume 2031 of *Lecture Notes in Computer Science*, pages 113–127, Genova, Italy, 2001. Springer-Verlag, Berlin.