

Invariant Generation through Strategy Iteration in Succinctly Represented Control Flow Graphs

THOMAS MARTIN GAWLITZA^a AND DAVID MONNIAUX^b

^a School of Information Technologies, The University of Sydney, Australia
e-mail address: gawlitza@it.usyd.edu.au

^b CNRS / VERIMAG Laboratory, Centre Équation, 2 avenue de Vignate, 38610 Gières, France
e-mail address: David.Monniaux@imag.fr

ABSTRACT. We consider the problem of computing numerical invariants of programs, for instance bounds on the values of numerical program variables. More specifically, we study the problem of performing static analysis by abstract interpretation using template linear constraint domains. Such invariants can be obtained by Kleene iterations that are, in order to guarantee termination, accelerated by widening operators. In many cases, however, applying this form of extrapolation leads to invariants that are weaker than the strongest inductive invariant that can be expressed within the abstract domain in use. Another well-known source of imprecision of traditional abstract interpretation techniques stems from their use of join operators at merge nodes in the control flow graph. The mentioned weaknesses may prevent these methods from proving safety properties.


The technique we develop in this article addresses both of these issues: contrary to Kleene iterations accelerated by widening operators, it is guaranteed to yield the strongest inductive invariant that can be expressed within the template linear constraint domain in use. It also eschews join operators by distinguishing all paths of loop-free code segments. Formally speaking, our technique computes the least fixpoint within a given template linear constraint domain of a transition relation that is succinctly expressed as an existentially quantified linear real arithmetic formula.

In contrast to previously published techniques that rely on quantifier elimination, our algorithm is proved to have optimal complexity: we prove that the decision problem associated with our fixpoint problem is Π_2^P -complete. Our procedure mimics a Π_2^P search.

1998 ACM Subject Classification: D.2.4, F.3.1 (D.2.1, D.2.4, D.3.1, E.1), F.3.2 (D.3.1).

Key words and phrases: static program analysis, abstract interpretation, fixpoint equation systems, strategy improvement algorithms, SMT solving.

* An earlier and shorter version of this article appeared in the proceedings of the 21st European Symposium on Programming (ESOP) 2011 [25].

^a This work was partially funded by the ANR project “ASOPT” 

^b VERIMAG is a joint laboratory of CNRS, Université Joseph Fourier and Grenoble INP.

1. INTRODUCTION

Static program analysis aims at deriving properties that are valid for all possible executions of a program, through an algorithmic processing of its source or object code. Examples of interesting properties include: “the program always terminates”; “the program never executes a division by zero”; “the program never dereferences a null pointer”; “the value of variable x always lies between 1 and 3”; “the output of the program is well-formed XHTML”. There is considerable practical interest in being able to prove such properties automatically, in particular for software used in safety-critical applications, e.g., in fly-by-wire flight control systems in aircraft [60].

1.1. Abstract interpretation. It is well-known that fully automatic, sound and complete program analysis is impossible for any nontrivial property regarding the final output of a program.¹ All analysis methods therefore suffer from at least one of the following limitations: they may be limited to programs with finite (and not too large) memory, or to bounded execution times; they may be *unsound* (they may infer untrue properties); or they may be *incomplete* (they fail to prove certain true properties). In this article, we use the *abstract interpretation framework* of Cousot and Cousot [18] to construct a static analysis technique that is sound, but incomplete.

Static analysis by abstract interpretation replaces the computation over concrete *reachable states* by computations over symbolically represented sets of concrete states. The sets are taken from an *abstract domain*. For instance, one may aim at computing, for each program point p and each program variable x , an interval in which the value of x is guaranteed to lie whenever the program reaches program point p . An analysis solely based on such intervals is known as *interval analysis* [17]. More refined numerical analyses include, for instance, finding for each program point an enclosing polyhedron for the vector of program variables [19]. By restricting the analysis to handle only sets found within a particular abstract domain (e.g., Cartesian products of intervals or convex polyhedra), one can make the problem tractable, at the expense of *over-approximation*. For instance, if the domain in use consists of convex shapes, only, non-convex invariants will necessarily get over-approximated.

In addition to the abstract domain not being able to represent the required properties, a major source of imprecision is the use of *widening operators* to enforce the convergence of Kleene iterations within finitely many iteration steps [18]. These operators extrapolate the first iterates of the Kleene sequence, say, of the intervals $[0, 1]$, $[0, 2]$, $[0, 3]$, \dots to a plausible limit, say $[0, +\infty)$, ensuring termination of the accelerated iteration. However, such an accelerated iteration may overshoot the target, leading to further over-approximations of the desired result. In order to regain precision lost by widening, one can then apply *narrowing*. In its simplest form, narrowing is a descending iteration towards a fixpoint that strengthens the invariant step by step. For more detailed information on Kleene iteration techniques in the context of abstract interpretation, we refer the reader to Cousot and Cousot [18]. Many variants of this basic iteration scheme have been proposed to alleviate the over-approximations introduced by widening [34, 35, 38]. However, all these techniques do not guarantee to find the strongest inductive invariant that can be expressed in the abstract domain in use.

¹This result, formally given within the framework of recursive function theory, is known as Rice’s theorem [54, p. 34][52, corollary B].

Let us illustrate the above mentioned weaknesses on the following simple example:

```
i = 0;
while (true) {
  if (i < 10) i = i+2;
  else goto _end; }
_end: printf("i = %d\n");
```

The strongest invariant, that is, the set of reachable states, is given by the proposition $i \in \{0, 2, 4, 6, 8, 10\}$, which, together with the exit condition $i \geq 10$, yields $i = 10$ as the only possible final value of i at program point `_end`. Interval analysis by Kleene iterations with widenings computes the intervals $[0, 0]$, $[0, 2]$, $[0, 4]$ and may then widen to $[0, +\infty)$. The narrowing phase yields the inductive invariant $i \in [0, 11]$. From this we can conclude that the final value of i is in the interval $[10, 11]$. The obtained interval $[0, 11]$ represents the strongest *inductive* invariant that can be expressed as an interval.² It is, however, not the strongest invariant expressible as an interval, which is $i \in [0, 10]$. The invariant $i \in [0, 10]$ is not inductive, because a state with $i = 9$ is mapped to a state with $i = 11$ by one iteration of the loop.

Unfortunately, small changes to the above program can make the widening/narrowing approach fail to produce a good invariant. Consider, for instance, the introduction of an additional non-deterministic choice, represented by the function `choice()`:

```
i = 0;
while (true) {
  if (choice()) {
    if (i < 10) i = i+2;
    else goto _end; } }
_end: printf("i = %d\n");
```

The program still outputs the value 10, whenever it terminates. The only difference from the first version of the program is that there is, in each iteration, a non-deterministic choice whether or not the original loop body is to be executed. If we perform the widening/narrowing technique on the modified version, the widening phase will produce the same result $[0, +\infty)$. However, the narrowing phase is now not able to regain any precision lost due to widening. The loop body represents the relation $\tau = \{(i, i) \mid i \in \mathbb{Z}\} \cup \{(i, i + 2) \mid i \in \mathbb{Z} \text{ and } i < 10\}$. This relation is reflexive, that is, $(i, i) \in \tau$ for all $i \in \mathbb{Z}$. The problem is of a general nature: Whenever the transition relation τ of a loop is reflexive, descending iterations fail to improve the inductive invariant obtained by widening.

Of course, on such a simple example, one could use simple tricks to get rid of the imprecision and recover the interval $[0, 11]$: remove the identity from the transition relation (this does not change the set of all (inductive) invariants), or try a form of widening with

²Some presentations of Hoare logic or static analysis call “invariant” what we refer to in this article as “inductive invariant”: a set (or a logical formula defining such a set) containing all initial states and stable by the transition relation. In our terminology, an invariant is merely a property true at all times. With these definitions, an inductive invariant is an invariant by induction on the length of the execution trace, thus the terminology; however an invariant is not necessarily inductive. Consider the initial state $(x, y) = (1, 0)$ and a transition consisting in a 45° clockwise rotation around $(0, 0)$: $(x, y) \in [-1, 1] \times [-1, 1]$ is an invariant (it is always true), but it is not inductive because $[-1, 1] \times [-1, 1]$ is not stable by this rotation.

thresholds, also known as widening “up to” [39]. However, such approaches are brittle and may fail for more complex programs.

1.2. Alternatives to the widening/narrowing approach. Because of the known weaknesses of the widening/narrowing approach, alternative methods have been proposed. Finding an inductive invariant in an abstract domain can be recast as solving a constraint system. Finding *the strongest* inductive invariant is then the problem of finding a *minimal* solution to the constraint system. The technique described in this article is related to two recently proposed approaches, which we shall now briefly describe.

Quantifier elimination. Monniaux [48] considers abstract domains where elements are defined by a logical formula I (more specifically, a conjunction of linear inequalities) that links the program variables to some parameters. For instance, intervals on two variables x, y are defined by $I := -l_x \leq x \leq u_x \wedge -l_y \leq y \leq u_y$, where l_x, u_x, l_y, u_y are the parameters. An element from the abstract domain defined by the template I is specified by an assignment of values to the parameters.

Consider a set of initial states given by a formula ι (in the above example, with free variables $\sigma = (x, y)$) and a transition relation given by τ (in the above example, with free variables $(\sigma, \sigma') = (x, y, x', y')$). I defines an inductive invariant for ι and τ if and only if

$$\forall \sigma . \iota(\sigma) \Rightarrow I(\sigma) \wedge \forall \sigma, \sigma' . (I(\sigma) \wedge \tau(\sigma, \sigma') \Rightarrow I(\sigma')). \quad (1.1)$$

Here, $I(\sigma)$ is the formula I as above and $I(\sigma')$ is the formula I with σ replaced by σ' . The free variables of formula (1.1) are the parameters in I . In the above example, they are l_x, u_x, l_y, u_y . Any satisfying assignment to these variables defines an inductive invariant from the abstract domain. A *least* inductive invariant in the abstract domain is then defined by constructing, using formula (1.1) as a building block, a formula whose solution is the minimal solution of (1.1), using that, for any formula F , $x_0 = \min\{x \mid F(x)\}$ if and only if

$$F(x_0) \wedge \forall x . (F(x) \Rightarrow x_0 \leq x). \quad (1.2)$$

The static analyzer then proceeds as follows: transform the loop into a set of initial states ι and a transition relation τ . From these formulas, construct Formula 1.2. Then, call a solver capable of dealing with quantified formulas, e.g, a quantifier elimination procedure or a lazy version thereof such as the one developed by Monniaux [47].

As an extension to this framework, ι and τ may have additional variables, e.g., precondition or system parameters. The formula defining the least inductive invariant will then take the invariant parameters as a partial function (in the mathematical sense, that is, as a binary related each input to at most one output) of these precondition or system parameters. By quantifier elimination and further processing of the formula, it is possible to turn this formula into a closed-form function, and even into executable code computing that function (a tree of if-then-else statements with assignments at the leaves).

This approach allows to effectively synthesize best abstract transformers ($\alpha \circ \tau \circ \gamma$ in the notation of Cousot and Cousot [18]). Unfortunately, quantifier elimination over linear real arithmetic is still very costly, despite the various recent works on this problem, and quantifier elimination over linear integer arithmetic and polynomial real arithmetic are even costlier.

The technique described in this article considers the same problem as the quantifier elimination approach, but without preconditions or system parameters. Our technique also uses a different algorithmic approach, called *max-strategy iteration*.

Strategy Iteration. In this article, we introduce a refinement of the *max-strategy iteration* technique of Gawlitza and Seidl [26] for template linear constraint domains. The phrase “strategy iteration”, also known as “policy iteration”, comes from game theory. Let us consider two-players zero-sum games: the outcome of such a game is a real number, the two players (the maximizer and the minimizer) aim at maximizing (respectively, minimizing) the outcome. Strategy iteration is a method for computing the optimal strategy for one of the players. It successively improves a strategy through the following two steps until an optimal strategy is found: (*Evaluation*) Evaluate the currently selected strategy; and (*Improvement*) try to improve the currently selected strategy w.r.t. the result of the evaluation.

The max-strategy iteration technique of Gawlitza and Seidl [26] for finding invariants is inspired by this game-theoretic approach. Instantiated on template linear constraint domains, it computes the strongest inductive invariant that can be represented by polyhedra of the form $P(b) = \{x \in \mathbb{R}^n \mid Tx \leq b\}$, where $T \in \mathbb{R}^{m \times n}$ is a template constraint matrix, which is fixed before the analysis is run (heuristics for finding a suitable matrix are out-of-scope for this article). The variable x is the vector of program variables. The template constraint matrix T is the counterpart of the template I from the quantifier elimination technique of Monniaux [48]. Given T , every vector $b \in \overline{\mathbb{R}}^m$ uniquely determines a polyhedron $P(b)$. The vector b contains the bounds on the linear functions that are represented by the rows of T . With the appropriate choice of T we can, among others, express the popular interval [17] and octagon [45, 46] abstract domains.

Similarly to Kleene iterations, the max-strategy improvement algorithm produces an ascending sequence of pre-fixpoints that are less than or equal to the least inductive invariant we are aiming for. The pre-fixpoints are obtained through convex optimization techniques, e.g., linear programming. In contrast to Kleene iterations, though, the algorithm converges to the least inductive invariant after at most exponentially many steps. Our conjecture is that it usually converges fast in practice, though one can concoct artificial examples that exhibit exponential behavior.

Trace partitioning. Max-strategy iteration rids us of imprecisions introduced by widening, but, per se, does not remove imprecisions introduced by another operation: the merging of information from different program paths at join nodes in the control flow graph. In this article, we introduce a refinement of max-strategy iteration where we distinguish the various execution paths, in a manner similar to the work of Monniaux [48], and Monniaux and Gonnord [49].

In most systems for static analysis by abstract interpretation, joins in the control-flow graph result in computations of least upper bounds in the abstract domain. For instance, consider abstract interpretation over general convex polyhedra on the following program:

```
if (x >= 0) y = x;
else y = -x;
if (y >= 1) z = 3.5/x;
```

The program divides 3.5 by the value of x provided that the absolute value of x is at least 1. A static analyzer that uses convex polyhedra as abstract domain may work as follows.




Figure 1: On the left: the graph of $y = |x|$ is the union of two half-lines, but computing their convex hull yields the grayed shape. By intersection with $y \geq 1$, we obtain the shape on the right, which contains points with $x = 0$ even though $y = |x| \wedge y \geq 1$ has no solution with $x = 0$.




Figure 2: Instead of considering two transitions (corresponding to a first if-then-else) followed by convex hull followed by two transitions (corresponding to a second if-then-else), as on the left, we get better precision by considering the four product transitions, as on the right.

After the first if-then-else statement, a convex hull is computed between the $x \geq 0 \wedge y = x$ and $x < 0 \wedge y = -x$ half-lines, resulting in a much larger polyhedron (see Fig. 1). The imprecision introduced by this operation prevents the analyzer from proving that a division by zero at line 3 is impossible.

One solution is to get rid of all convex hulls corresponding to control flow joins by removing all control flow joins, except those corresponding to loop headers, by combining control flow edges. For instance, n successive if-then-else constructs can be turned into an expanded system of 2^n transitions (Figure 2 shows this construction for $n = 2$). This is close to the *trace partitioning* approach of Rival and Mauborgne [53].³ One could therefore run this exponential transformation first, and then run max-strategy iteration or min-strategy iteration (Sec. 1.4). However, this transformation causes an exponential blowup and is therefore clearly not scalable.

³Trace partitioning analyses each program statement in different contexts according to an abstraction of the history of the control trace; thus, if a statement is preceded by n tests, it can potentially analyze this statement in 2^n contexts. Because of this exponential blowup of maximal partitioning, trace partitioning techniques, including those implemented in Astrée [8, 9], use heuristics to “fold” abstract elements together using join operations.

In this article, we describe an algorithm that yields the same result as max-strategy iteration on this exponentially larger system. Our algorithm uses only polynomial space. It achieves this by keeping the exponentially large system implicit.

Path focusing. Henry et al. [40], Monniaux and Gonnord [49] propose to run the classical Kleene iterations with widening and narrowing scheme not on the original control-flow graph, but on this exponentially larger system. In this approach, iterations are run on a distinguished subset of the original control nodes, such that all cycles in the original control flow graphs cross at least one of these distinguished nodes, using transitions corresponding to the simple paths between these distinguished nodes in the original control flow graph. The expanded control multigraph is kept implicit: the transitions, corresponding to simple paths in the original graph, are obtained on demand as solutions to SMT problems. This approach has the following advantages:

- (1) It fully does away with imprecisions introduced by “join” operations, except those corresponding to loops.
- (2) The transition relations on the simple paths may be accelerable. That is, they can be dealt with through acceleration techniques (cf. Sec. 1.4, [32, 33, 43]).
- (3) While it uses widening operators, it does away with some of the imprecisions they introduce by focusing on one path at a time, which allows the use of narrowing iterations even on programs where they fail to yield better precision with the classical iteration scheme.

The technique we present in this article combines the idea of implicit representation with max-strategy iteration.

1.3. Contributions. The main contribution of this article is an algorithm that computes the strongest inductive invariant of the expanded transition system (which allows higher precision for abstract interpretation) without actually constructing it. We shall see later the exact definition, but here is an interesting particular case (the general result allows more complex control flow): given a $m \times n$ matrix A , an initial value $\iota \in \mathbb{Q}^n$ and a transition relation τ over \mathbb{Q}^n , defined by a formula over variables $x_1, \dots, x_n, x'_1, \dots, x'_n$, built with non-strict linear (in)equalities, \wedge , \vee and prenex \exists , compute the least set of the form $P(b) = \{x \in \mathbb{R}^n \mid Ax \leq b\}$ (that is, compute b) containing ι and stable by the transition relation τ ; equivalently, find the least loop invariant of the form $Ax \leq b$ for the loop with initial state ι and loop body expressed by τ .

Our algorithm can be performed in polynomial space and exponential time. It works in a demand-driven fashion: elements from the exponentially-sized sets of strategies and loop-free paths are enumerated only as needed, and one can thus hope that they will not all be enumerated, which seems to be confirmed by our preliminary experiments.

We also consider the following associated decision problem, which we shall later make more formal:

“Given a control-flow graph (with N vertices) and transition relations written as existentially quantified first-order linear real arithmetic formulas, a family A_1, \dots, A_N of matrices, an initial control state and a “bad” control state b , does there exist vectors b_1, \dots, b_N such that $A_1x \leq b_1 \wedge \dots \wedge A_Nx \leq b_N$ forms an inductive invariant proving that b is unreachable?”

We show this problem to be Σ_2^P -complete (at the second level of the polynomial time hierarchy [50, ch. 17]), even if $N = 1$ and the matrix is 1×1 . Equivalently, the negated problem (abstract reachability of a statement) is shown to be Π_2^P -complete. Assuming the polynomial hierarchy does not collapse, this means that this problem can be solved in polynomial space, but is harder than **NP**-complete and **coNP**-complete problems. This clearly justifies the use of an exponential-time algorithm.

1.4. Other related Work. Many approaches have been proposed to address the imprecisions caused by widening operators. We now briefly describe approaches related to ours, in addition to those that we directly build upon (Sec. 1.2). Halbwachs et al. [39] proposed widening “up to” (an idea resurrected in the Astrée system as widening with thresholds [8, 9]), which extracts syntactic hints for limiting widening. Bagnara et al. [4, 5] proposed improvements over the “classical” widenings on linear constraint domains [37]. Gopan and Reps [34] introduced “look-ahead widening” [34] and “guided iterations” [35]: standard widening-based analysis is applied to a sequence of syntactic restrictions of the original program, which ultimately converges to the whole program; the idea is to distinguish phases or modes of operation in order to make the widening more precise. Some other techniques fully do away with widenings [13, 15, 55], for instance by expressing the invariants as solutions of a mathematical programming problem [36], and thus the least invariant in the domain as an optimal solution to this problem.

In some cases, it is possible to compute exactly the transitive closure of the transition relation, or the application of the transitive closure to given initial states, or at least to compute a good over-approximation thereof. Such *acceleration* techniques [32, 33, 43] tend to have difficulties dealing with programs where the control flow is not flat (multiple paths within the loop body).

In Section 1.2, we sketched max-strategy iteration by an analogy to solving games where “max” operations correspond to control-flow joins and “min” operations to guards. If instead of choosing arguments to “max” operators, the strategy chooses them for “min” operators, we obtain min-strategy iteration [14, 24]. Min-strategy iteration solves a sequence of fixpoint problems with decreasing values always weaker or equivalent to the strongest inductive invariant in the domain. In general, this sequence does not necessarily converge to this least inductive invariant, but it does so under certain conditions (e.g., when all abstract transformers are non-expansive [1]). We investigated applying our “implicit representation” idea to the min-strategy approach, but encountered a stumbling block: while it is possible to decide whether a max-strategy is improvable using SMT solving on quantifier-free formulas, the equivalent for min-strategies necessitated quantified formulas, which defeats the purpose of doing away with quantifier elimination techniques.

2. BASICS

2.1. Notations. $\mathbb{B} = \{0, 1\}$ denotes the set of Boolean values. The set of real numbers (resp. the set of rational numbers) is denoted by \mathbb{R} (resp. \mathbb{Q}). The complete linearly ordered set $\mathbb{R} \cup \{-\infty, \infty\}$ is denoted by $\overline{\mathbb{R}}$, similarly $\mathbb{Q} \cup \{-\infty, \infty\}$ is denoted by $\overline{\mathbb{Q}}$. For any expression (resp. term) e , we write $e[e_1/\mathbf{x}_1, \dots, e_k/\mathbf{x}_k]$ to denote the expression (resp. term) that is obtained from e by simultaneously replacing all occurrences of the variables $\mathbf{x}_1, \dots, \mathbf{x}_k$ by e_1, \dots, e_k .

A partially ordered set \mathbb{D} is called a *lattice* if and only if any two elements $x, y \in \mathbb{D}$ have a greatest lower bound and a least upper bound, denoted respectively by $x \wedge y$ and $x \vee y$. It is a *complete lattice* if and only if any subset $X \subseteq \mathbb{D}$ has a greatest lower bound and a least upper bound, denoted by $\bigwedge X$ and $\bigvee X$. The least element $\bigvee \emptyset$ of a complete lattice is denoted by \perp . The greatest element $\bigwedge \emptyset$ is denoted by \top .

Assume that \mathbb{D}_1 and \mathbb{D}_2 are partially ordered by \leq_1 and \leq_2 , respectively. A function $f : \mathbb{D}_1 \rightarrow \mathbb{D}_2$ is called *monotone* if and only if $f(x) \leq_2 f(y)$ for all $x, y \in \mathbb{D}_1$ with $x \leq_1 y$. We shall often use the following fundamental result:

Theorem 2.1 (Knaster/Tarski [62]). *Let \mathbb{D} be a complete lattice and $f : \mathbb{D} \rightarrow \mathbb{D}$ monotone. The operator f has a least fixpoint and a greatest fixpoint, respectively denoted by μf and νf . Moreover, we have $\mu f = \bigwedge \{x \in \mathbb{D} \mid f(x) \leq x\}$ and $\nu f = \bigvee \{x \in \mathbb{D} \mid x \leq f(x)\}$. \square*

We denote the transpose of a matrix A by A^\top . For $x \in \overline{\mathbb{R}}$, we denote the column vector $(x, \dots, x)^\top$ by \underline{x} . We denote the i -th row (resp. the j -th column) of a matrix A by A_i . (resp. $A_{\cdot j}$). Accordingly, $A_{i,j}$ denotes the entry in the i -th row and the j -th column. We also use this notation for vectors and mappings $f : X \rightarrow Y^k$, i.e., for all $i \in \{1, \dots, k\}$, the mapping $f_i : X \rightarrow Y$ is given by $f_i(x) = (f(x))_i$ for all $x \in X$. The set $\overline{\mathbb{R}}^n$ is partially ordered by the component-wise extension of \leq , which we again denote by \leq . That is, for all $x, y \in \overline{\mathbb{R}}^n$, $x \leq y$ if and only if $x_i \leq y_i$ for all $i \in \{1, \dots, n\}$.

A mapping $f : \overline{\mathbb{R}}^n \rightarrow \overline{\mathbb{R}}^m$ is called *affine* if and only if there exist $A \in \mathbb{R}^{m \times n}$ and $b \in \overline{\mathbb{R}}^m$ such that $f(x) = Ax + b$ for all $x \in \overline{\mathbb{R}}^n$. Here, we use the convention $-\infty + \infty = -\infty$. Observe that f is monotone if all entries of A are non-negative. A mapping $f : \overline{\mathbb{R}}^n \rightarrow \overline{\mathbb{R}}$ is called *weak-affine* if and only if there exist $a \in \mathbb{R}^n$ and $b \in \overline{\mathbb{R}}$ such that $f(x) = a^\top x + b$ for all $x \in \overline{\mathbb{R}}^n$ with $f(x) \neq -\infty$. A mapping $f : \overline{\mathbb{R}}^n \rightarrow \overline{\mathbb{R}}^m$ is called *weak-affine* if and only if there exist weak-affine mappings $f_1, \dots, f_m : \overline{\mathbb{R}}^n \rightarrow \overline{\mathbb{R}}$ such that $f = (f_1, \dots, f_m)^\top$. Every affine mapping is weak-affine, but not vice-versa. In this article, we are concerned with mappings that are point-wise minimums of finitely many monotone and weak-affine mappings. Note that these mappings are in particular concave, i.e., the set of points below the graph of the function is convex.

2.2. Linear Programming. Linear programming aims at optimizing a linear objective function with respect to linear constraints. In this article, we consider linear programming problems (LP problems for short) of the form $\sup \{c^\top x \mid x \in \mathbb{R}^n, Ax \leq b\}$. Here, $A \in \mathbb{R}^{m \times n}$, $b \in \mathbb{R}^m$, and $c \in \mathbb{R}^n$ are the inputs. The convex closed polyhedron $\{x \in \mathbb{R}^n \mid Ax \leq b\}$ is called the *feasible space*. The LP problem is called *infeasible* if and only if the feasible space is empty. An element of the feasible space, is called *feasible solution*. A feasible solution x that maximizes $c^\top x$ is called *optimal solution*.

If A and b consist of rational entries, only, then the feasible space is nonempty if and only if it contains a rational point. An optimal solution exists if and only if there exists a rational one. In this article, we always assume that all numbers in the input are rational.

LP problems can be solved in polynomial time through the ellipsoid method [41] and interior point methods [57]. However, the running-time of these algorithms crucially depends on the sizes of occurring numbers. At the danger of an exponential running-time in contrived cases, we can also instead rely on the simplex algorithm: its worst-case running-time does not depend on the sizes of occurring numbers (given that arithmetic operations, comparison, storage and retrieval for numbers are counted for $\mathcal{O}(1)$). See for example Dantzig [20], Schrijver [57] for more information on linear programming.

2.3. SAT modulo linear real arithmetic. The set of SAT modulo linear real arithmetic formulas Φ is defined through the following grammar:

$$e ::= c \mid x \mid e_1 + e_2 \mid c \cdot e' \quad \Phi ::= a \mid e_1 \leq e_2 \mid e_1 < e_2 \mid \Phi_1 \vee \Phi_2 \mid \Phi_1 \wedge \Phi_2 \mid \neg\Phi' \quad (2.1)$$

Here, $c \in \mathbb{Q}$ is a constant, x is a real valued variable, e, e', e_1, e_2 are real-valued linear expressions, a is a Boolean variable and $\Phi, \Phi', \Phi_1, \Phi_2$ are formulas. An *interpretation* I for a formula Φ is a mapping that assigns a real value to every real-valued variable and a Boolean value to every Boolean variable. We write $I \models \Phi$ for “ I is a *model* of Φ ”. That is, we firstly inductively define a function $\llbracket e \rrbracket$ that evaluates a linear expression e as follows:

$$\llbracket c \rrbracket I = c \quad \llbracket x \rrbracket I = I(x) \quad \llbracket e_1 + e_2 \rrbracket I = \llbracket e_1 \rrbracket I + \llbracket e_2 \rrbracket I \quad \llbracket c \cdot e' \rrbracket I = c \cdot \llbracket e' \rrbracket I \quad (2.2)$$

Secondly, we inductively define the relation \models as follows:

$$\begin{aligned} I \models a &\iff I(a) = 1 & I \models e_1 \leq e_2 &\iff \llbracket e_1 \rrbracket I \leq \llbracket e_2 \rrbracket I \\ I \models e_1 < e_2 &\iff \llbracket e_1 \rrbracket I < \llbracket e_2 \rrbracket I & I \models \Phi_1 \vee \Phi_2 &\iff I \models \Phi_1 \text{ or } I \models \Phi_2 \\ I \models \Phi_1 \wedge \Phi_2 &\iff I \models \Phi_1 \text{ and } I \models \Phi_2 & I \models \neg\Phi' &\iff I \not\models \Phi' \end{aligned} \quad (2.3)$$

A formula is called *satisfiable* if and only if it has at least one model. A formula has a model if and only if it has a rational model.

The problem of deciding the satisfiability of SAT modulo linear real arithmetic formulas is **NP**-complete. There nevertheless exist efficient solver implementations for this decision problem, generally based on the DPLL(T) approach, an extension of the DPLL algorithm for SAT to richer logics. For more information see for example Biere et al. [7], Dutertre and de Moura [22], and Kroening and Strichman [42]. Such implementations, on satisfiable instances, can provide a model over Booleans and rational numbers.

In order to simplify notations we also allow matrices, vectors, the relations $\geq, >, \neq, =$, and the Boolean constants 0 and 1 to occur in SAT modulo linear real arithmetic formulas.

3. THE FRAMEWORK

3.1. Control Flow Graphs and Collecting Semantics. In this article, we model programs as control flow graphs, i.e., a *program* G is a triple (N, E, \mathbf{st}) , where

- (1) N is a finite set of *program points*,
- (2) $E \subseteq N \times \mathbf{Stmt} \times N$ is a finite set of *control-flow edges*, and
- (3) $\mathbf{st} \in N$ is the *start program point*.

A program uses n real-valued variables $\mathbf{x}_1, \dots, \mathbf{x}_n$. A state is described by a vector $x \in \mathbb{R}^n$. We assign a *collecting semantics* $\llbracket s \rrbracket : 2^{\mathbb{R}^n} \rightarrow 2^{\mathbb{R}^n}$ to each statement $s \in \mathbf{Stmt}$. The collecting semantics $\llbracket s \rrbracket$ is an operator that assigns a set $\llbracket s \rrbracket(X)$ of possible states after the execution of s to a set X of possible states before the execution of s . The set \mathbf{Stmt} of statements is specified subsequently. The *collecting semantics* V of a program $G = (N, E, \mathbf{st})$ is finally defined as the least solution of the following constraint system:

$$\mathbf{V}[\mathbf{st}] \supseteq \mathbb{R}^n \quad \mathbf{V}[v] \supseteq \llbracket s \rrbracket(\mathbf{V}[u]) \quad \text{for all } (u, s, v) \in E. \quad (3.1)$$

Here, for any $v \in N$, the variable $\mathbf{V}[v]$ takes values in $2^{\mathbb{R}^n}$. The components of the collecting semantics V are denoted by $V[v]$ for all $v \in N$. Throughout this article, we will usually denote variables in bold face, and values in normal face.

3.2. Statements. The set **Stmt** of all statements is the set of all SAT modulo linear real arithmetic formulas without Boolean variables and without negation. Note that non-strict and strict inequality constraints are permitted. The formula $e_1 \neq e_2$ is also permitted, since it is an abbreviation for $e_1 < e_2 \vee e_2 < e_1$. We can (in linear time) transform any SAT modulo linear real arithmetic formula without Boolean variables into this form by pushing negations to the leaves.

The \mathbb{R} -valued variables $\mathbf{x}_1, \dots, \mathbf{x}_n$ and $\mathbf{x}'_1, \dots, \mathbf{x}'_n$, that may occur in the formula, play a particular role. The values of the variables $\mathbf{x}_1, \dots, \mathbf{x}_n$ represent the values of the program variables before executing the statement, and the values of the variables $\mathbf{x}'_1, \dots, \mathbf{x}'_n$ represent the values of the program variables after executing the statement. For convenience, we denote the vectors $(\mathbf{x}_1, \dots, \mathbf{x}_n)^\top$ and $(\mathbf{x}'_1, \dots, \mathbf{x}'_n)^\top$ also by \mathbf{x} and \mathbf{x}' , respectively. In addition to $\mathbf{x}_1, \dots, \mathbf{x}_n$ and $\mathbf{x}'_1, \dots, \mathbf{x}'_n$, the statement may also include other variables, which may stand for intermediate values computed (or non-deterministically chosen) during the execution of a program statement. Conceptually, these variables are existentially quantified.

We could also add Boolean variables, at the expense of some additional complexity in definitions, theorems and proofs. Note that this would not increase the expressiveness, since a Boolean variable \mathbf{y} can be simulated by a real variable $\tilde{\mathbf{y}}$ by replacing all occurrences of \mathbf{y} by $\tilde{\mathbf{y}} = 1$, all occurrences of $\neg \mathbf{y}$ by $\tilde{\mathbf{y}} = 0$, and conjoining $(\tilde{\mathbf{y}} = 0 \vee \tilde{\mathbf{y}} = 1)$ to the formula. In practice, the direct support of Boolean variables may be beneficial for the efficiency. More generally, we can accommodate any formula feature that just expresses disjunctions in a compact way; the only requirement is not to generate negations.

The *collecting semantics* $\llbracket s \rrbracket : 2^{\mathbb{R}^n} \rightarrow 2^{\mathbb{R}^n}$ of a statement $s \in \mathbf{Stmt}$ is defined by

$$\llbracket s \rrbracket(X) := \{x' \in \mathbb{R}^n \mid \exists x \in X. s[x/\mathbf{x}, x'/\mathbf{x}'] \text{ is satisfiable}\} \quad \text{for all } X \subseteq \mathbb{R}^n. \quad (3.2)$$

Consider the following C-code snippet:

```

if (x_1 >= 0)
  x_2 = x_1;
else
  x_2 = -x_1;

```

Assume that `x.1` and `x.2` are of type **int** and that they are the only numerical variables. The effect of the C code snippet can be abstracted by the statement

$$\mathbf{x}'_1 = \mathbf{x}_1 \wedge ((\mathbf{x}_1 \geq 0 \wedge \mathbf{x}'_2 = \mathbf{x}_1) \vee (\mathbf{x}_1 < 0 \wedge \mathbf{x}'_2 = -\mathbf{x}_1)) \quad (3.3)$$

Note that a conjunct $\mathbf{x}'_i = \mathbf{x}_i$ is needed for all variables that do not change their values.

A statement s is called *merge-simple* if and only if it is in disjunctive normal form, i.e., s is of the form $s_1 \vee \dots \vee s_k$, where the statements s_1, \dots, s_k do not use the Boolean connector \vee . Any statement can be rewritten into an equivalent merge-simple statement in exponential time and space using distributivity. The crux of our main result is that our algorithm never needs to compute such an exponentially-sized disjunctive normal form.

If we convert Statement (3.3) into an equivalent merge-simple statement using distributivity, we get:

$$(\mathbf{x}'_1 = \mathbf{x}_1 \wedge \mathbf{x}_1 \geq 0 \wedge \mathbf{x}'_2 = \mathbf{x}_1) \vee (\mathbf{x}'_1 = \mathbf{x}_1 \wedge \mathbf{x}_1 < 0 \wedge \mathbf{x}'_2 = -\mathbf{x}_1) \quad (3.4)$$

A merge-simple statement s that does not use the Boolean connector \vee at all is called *sequential*. Intuitively, sequential statements correspond to straight-line sequences of basic

blocks. The merge-simple statement (3.4) non-deterministically chooses between executing one of the following sequential statement:

$$\mathbf{x}'_1 = \mathbf{x}_1 \wedge \mathbf{x}_1 \geq 0 \wedge \mathbf{x}'_2 = \mathbf{x}_1 \qquad \mathbf{x}'_1 = \mathbf{x}_1 \wedge \mathbf{x}_1 < 0 \wedge \mathbf{x}'_2 = -\mathbf{x}_1 \qquad (3.5)$$

3.3. Abstract Semantics. Let \mathbb{D} be a complete lattice (for instance the complete lattice of all n -dimensional closed real intervals). Assume that $\alpha : 2^{\mathbb{R}^n} \rightarrow \mathbb{D}$ and $\gamma : \mathbb{D} \rightarrow 2^{\mathbb{R}^n}$ form a Galois connection, i.e., for all $X \subseteq \mathbb{R}^n$ and all $d \in \mathbb{D}$, $\alpha(X) \leq d$ if and only if $X \leq \gamma(d)$. The *abstract semantics* $\llbracket s \rrbracket^\sharp : \mathbb{D} \rightarrow \mathbb{D}$ of a statement s is then defined by

$$\llbracket s \rrbracket^\sharp := \alpha \circ \llbracket s \rrbracket \circ \gamma. \qquad (3.6)$$

Remark that we have chosen to use the best abstract transformer, i.e., the most precise abstract semantics. All that was needed for soundness is that $\llbracket s \rrbracket \circ \gamma(d) \subseteq \gamma \circ \llbracket s \rrbracket^\sharp(d)$ for all $d \in \mathbb{D}$. Our choice of $\llbracket s \rrbracket^\sharp(d)$, however, is the most accurate sound value.

The *abstract semantics* V^\sharp of a program $G = (N, E, \mathbf{st})$ is the least solution of the following constraint system:

$$\mathbf{V}^\sharp[\mathbf{st}] \geq \alpha(\mathbb{R}^n) \qquad \mathbf{V}^\sharp[v] \geq \llbracket s \rrbracket^\sharp(\mathbf{V}^\sharp[u]) \quad \text{for all } (u, s, v) \in E \qquad (3.7)$$

Here, for any $v \in N$, the variable $\mathbf{V}^\sharp[v]$ takes values in \mathbb{D} . The components of the abstract semantics V^\sharp are denoted by $V^\sharp[v]$ for all $v \in N$. The abstraction is sound, i.e., the abstract semantics V^\sharp safely over-approximates the collecting semantics V , i.e., $\gamma(V^\sharp[v]) \supseteq V[v]$ for all $v \in N$.

3.4. Template Linear Constraints. In this article we restrict our considerations to *template linear constraint domains* as introduced by Sankaranarayanan et al. [56]. We assume that a *template constraint matrix* $T \in \mathbb{R}^{m \times n}$ is given. For technical convenience, we always assume w.l.o.g. that $m \geq 1$ and each row of T contains at least one non-zero entry. The template linear constraint domain can be identified with the set $\overline{\mathbb{R}}^m$. As shown by Sankaranarayanan et al. [56], the abstraction $\alpha : 2^{\mathbb{R}^n} \rightarrow \overline{\mathbb{R}}^m$ and the concretization $\gamma : \overline{\mathbb{R}}^m \rightarrow 2^{\mathbb{R}^n}$, which are defined by

$$\gamma(d) := \{x \in \mathbb{R}^n \mid Tx \leq d\} \qquad \text{for all } d \in \overline{\mathbb{R}}^m, \text{ and} \qquad (3.8)$$

$$\alpha(X) := \bigwedge \{d \in \overline{\mathbb{R}}^m \mid \gamma(d) \supseteq X\} \qquad \text{for all } X \subseteq \mathbb{R}^n, \qquad (3.9)$$

form a Galois connection.

The template linear constraint domains contain *intervals*, *zones*, and *octagons* [45, 46], with appropriate choices of the template constraint matrix T [56]. For instance, if we have two variables x and y , and we abstract each variable by an interval as $x \in [-l_x, u_x]$ and $y \in [-l_y, u_y]$, the vector d is formed of (l_x, l_y, u_x, u_y) . Here, the matrix T is given by:

$$T = \begin{pmatrix} -1 & 0 \\ 0 & -1 \\ 1 & 0 \\ 0 & 1 \end{pmatrix}$$

and thus the concretization expresses:

$$\gamma \begin{pmatrix} l_x \\ l_y \\ u_x \\ u_y \end{pmatrix} = \left\{ \begin{pmatrix} x \\ y \end{pmatrix} \mid x \in [-l_x, u_x], y \in [-l_y, u_y] \right\} = \left\{ \begin{pmatrix} x \\ y \end{pmatrix} \mid \begin{pmatrix} -1 & 0 \\ 0 & -1 \\ 1 & 0 \\ 0 & 1 \end{pmatrix} \begin{pmatrix} x \\ y \end{pmatrix} \leq \begin{pmatrix} l_x \\ l_y \\ u_x \\ u_y \end{pmatrix} \right\}$$

While intervals, zones, and octagons are somewhat “obvious” choices, a common discussion with respect to template domains is how to find the templates, as opposed to the domain of convex polyhedra, where the convex hull and widening operations somewhat “discover” interesting directions in space. In this article, we shall assume that template matrices are given and refrain from discussing how they were obtained.

4. IMPROVING THE PRECISION OF THE ABSTRACTION


Most abstract interpretation techniques consider a control-flow graph with transitions expressed as sequential statements only (see formal definition in Sec. 3.2), that is, composed of atomic guards and assignments. An if-then-else construct with simple constructs (e.g., assignments) in both branches is thus expressed as two sequential statements, and a sequence of two such if-then-else constructs (one from point A to point B and one from B to C) is expressed as on the left of Figure 2: two sequential statements between A and B , and two sequential statements between B and C . As noted in the introduction (Sec. 1.2), abstract interpretation techniques usually abstract the set of reachable states at point B . This may result in spurious states being considered in the abstraction, which in turn may result in the analysis tool being unable to prove desirable properties.

In this article, we apply an idea that is very similar to the *path focusing* technique of Monniaux and Gonnord [49]. Given a program expressed as a control-flow graph with sequential statements on the edges, we first compute a *feedback vertex set* (a.k.a. *cut-set*) S , that is, a set of control nodes (the *feedback vertexes*) such that removing them cuts all cycles in the graph. Our original program is equivalent to a program where the only control nodes are those in the feedback vertex set, but edges carry arbitrary statements instead of sequential statements only (cf. Sec. 3.2). The results of program analyses on this new graph, at nodes from the feedback vertex set S , are sound invariants for the original program. If information is needed at other nodes, we can compute it from the information we have for the nodes from S .

Since methods for obtaining compact formulas expressing these statements from the original program have already been described in other publications [49], we do not explain them in detail. Instead, we provide an example.

Example 4.1 (Running Example). Throughout this article we use the following C-code snippet as a running example:

```
int x_1, x_2;
x_1 = 0;
while (x_1 <= 1000) {
    x_2 = -x_1;
    if (x_2 < 0) x_1 = -2 * x_1;
    else x_1 = -x_1 + 1; }
```

Figure 3: The program G_1 of the running example

This C-code snippet is abstracted through the program $G_1 = (N_1, E_1, \mathbf{st})$ depict in Figure 3. However, it is not necessary to apply abstraction at every program point, i.e., to assign an abstract value to each program point. It suffices to apply abstraction at a vertex feedback set of G_1 . Since all loops contain the program point 1, $\{1\}$ is a feedback vertex set of G_1 . Equivalently to applying abstraction only at program point 1, we can rewrite the control-flow graph G_1 into a control-flow graph $G = (N, E, \mathbf{st})$ that is equivalent w.r.t. the collecting semantic, but contains just the program point \mathbf{st} and the program points from the vertex feedback set $\{1\}$. The result of this transformation — the control-flow graph G — is shown in Figure 4(a) (Page 18).

The programs G_1 and G are equivalent w.r.t. their collecting semantics, i.e., $V[v] = V_1[v]$ for all $v \in N$. Here, V_1 denotes the collecting semantics of G_1 and V denotes the collecting semantics of G . W.r.t. to the abstract semantics, G is usually more precise than G_1 , because we reduced the number of merge points. In general, we only have $V^\sharp[v] \subseteq V_1^\sharp[v]$ for all $v \in N$, where V_1^\sharp denotes the abstract semantics of G_1 and V^\sharp denotes the abstract semantics of G . This is independent of the abstract domain.⁴ \square

Let us make a few last remarks regarding the feedback vertex set. Abstract interpretation techniques usually use such a set to select widening points [16, §4.1.2]. In contrast, our method uses this set to select the nodes where it over-approximates the set of reachable states; it does not over-approximate the set of reachable states at other nodes; widening is not involved at all. Finding a feedback vertex set of minimal cardinality is an NP-complete problem if the control-flow graph is arbitrary; such a set can however be found in linear time if the control-flow graph is *reducible* (in short, if loops have a single entry point) [59], which is the case for control-flow graphs directly obtained from structured programs (the method extends to certain irreducible graphs). The control-flow graph may however become irreducible if certain optimizations or partitioning techniques are used. A common heuristic is, for structured programs, to use loop headers, and for unstructured programs to use the targets of back edges from a depth-first traversal [10, 11]; this heuristic does not guarantee that the feedback vertex set is minimal with respect to inclusion ordering, let alone cardinality.

⁴We assume that we have given a Galois-connection and thus in particular monotone best abstract transformers.

5. BASIC OBSERVATIONS

We now note down basic properties of the abstract semantics.

5.1. Abstract Semantics of Statements. Our first observation is that, for all *sequential* statements s and all $d \in \overline{\mathbb{R}}^m$, $\llbracket s \rrbracket^\sharp(d)$ can be computed efficiently.

Lemma 5.1 (Sequential Statements). *Let s be a sequential statement and $d \in \overline{\mathbb{R}}^m$. The operator $\llbracket s \rrbracket^\sharp$ is a point-wise minimum of finitely many monotone and weak-affine operators. For all $d \in \overline{\mathbb{R}}^m$, $\llbracket s \rrbracket^\sharp(d)$ can be computed in polynomial time through linear programming.*

Proof. Let $i \in \{1, \dots, m\}$. We get:

$$\llbracket s \rrbracket_i^\sharp(d) = \sup \{T_i.x' \mid x' \in \llbracket s \rrbracket(\gamma(d))\} \quad (5.1)$$

$$= \sup \{T_i.x' \mid x' \in \mathbb{R}^n \text{ and } \exists x. Tx \leq d \text{ and } s[x/\mathbf{x}, x'/\mathbf{x}'] \text{ is satisfiable}\} \quad (5.2)$$

Equation 5.2 follows from Equation 5.1 by expansion of the concrete semantics $\llbracket s \rrbracket$ into a SMT-formula and of $\gamma(d)$ into $Tx \leq d$. Since s does not contain disjunctions, the optimization problem in (5.2) aims at optimizing a linear objective function w.r.t. linear constraints (equalities, strict inequalities, and non-strict inequalities). The optimal value of this optimization problem can be computed in polynomial time through linear programming. To check feasibility by standard linear programming techniques (which only allow non-strict inequalities), we can replace every strict inequality $e_1 < e_2$ by the non-strict inequality $e_1 \leq e_2 - \epsilon$, where ϵ is appropriately small. Such an appropriately small ϵ can be computed in polynomial time. Provided that the optimization problem is feasible, we can then replace $s[x/\mathbf{x}, x'/\mathbf{x}']$ by $s[x/\mathbf{x}, x'/\mathbf{x}'][\leq/\leq]$. Here, $s[\leq/\leq]$ denotes the statement obtained from s by replacing every strict inequality relation by a non-strict inequality relation. The optimal value of the obtained linear programming problem is equal to the optimal value of the optimization problem (5.2).

It remains to show that $\llbracket s \rrbracket_i^\sharp$ is a point-wise minimum of finitely many monotone and weak-affine operators. Since $s[x/\mathbf{x}, x'/\mathbf{x}'][\leq/\leq]$ is a conjunction of non-strict linear inequalities, there exist matrices A , A' and A'' and a vector b such that, for all x and x' , $s[x/\mathbf{x}, x'/\mathbf{x}'][\leq/\leq]$ is satisfiable if and only if there exists a x'' such that $Ax + A'x' + A''x'' \leq b$ (the vector x'' stands for the other variables in s , which are implicitly existentially quantified). Thus, the optimization problem (5.2) can be rewritten as follows:

$$\llbracket s \rrbracket_i^\sharp(d) = \sup \{T_i.x' \mid x' \in \mathbb{R}^n, \exists x \in \mathbb{R}^n. \exists x'' \in \mathbb{R}^q. Tx \leq d \text{ and } Ax + A'x' + A''x'' \leq b\} \quad (5.3)$$

Strong duality [12], also known as Farkas' lemma, thus gives us, provided that $\llbracket s \rrbracket_i^\sharp(d) > -\infty$, i.e., the optimization problem is feasible, the following equation:

$$\llbracket s \rrbracket_i^\sharp(d) = \inf \left\{ d^\top y_1 + b^\top y_2 \mid y_1, y_2 \geq 0, T^\top y_1 + A^\top y_2 = 0, A''^\top y_2 = 0, A'^\top y_2 = T_i^\top \right\} \quad (5.4)$$

Since $y_1 \geq 0$ for all feasible solutions of the linear programming problem in (5.4), $\llbracket s \rrbracket_i^\sharp$ coincides with a point-wise infimum of monotone and affine operators on the set $\{d \in \overline{\mathbb{R}}^m \mid \llbracket s \rrbracket_i^\sharp(d) > -\infty\}$. That is, $\llbracket s \rrbracket_i^\sharp$ is a point-wise infimum of monotone and weak-affine operators. Since the optimal value, provided that it exists, is attained at the vertices of the feasible space (finitely many), the point-wise infimum is a point-wise minimum of finitely many monotone and weak-affine operators. \square

The max-strategy improvement algorithm we adapt in this article heavily relies on the fact that, for all sequential statements s , $\llbracket s \rrbracket^\sharp$ is a point-wise minimum of finitely many monotone and weak-affine operators. The latter statement especially implies that $\llbracket s \rrbracket^\sharp$ is concave (see Gawlitza and Seidl [29] for precise definitions).

The number of vertices in the feasible space of the point-wise infimum in (5.4) may be exponential in the size of the original problem, and thus the representation as a point-wise minimum of finitely many monotone and weak-affine operators might contain an exponential number of such operators. This is not a problem since our algorithm never computes this decomposition explicitly.

Any polynomial-time method for evaluating the abstract semantics of sequential statements can be used to derive a polynomial-time method for evaluating merge-simple statements.

Lemma 5.2 (Merge-Simple Statements). *Let s be a merge-simple statement. The operator $\llbracket s \rrbracket^\sharp$ is a point-wise maximum of finitely many point-wise minima of finitely many monotone and weak-affine mappings. For all $d \in \mathbb{R}^m$, $\llbracket s \rrbracket^\sharp(d)$ can be computed in polynomial time through linear programming.*

Proof. Let $s \equiv s_1 \vee \dots \vee s_k$, where s_1, \dots, s_k are sequential statements. Since $\llbracket s \rrbracket^\sharp(d) = \llbracket s_1 \rrbracket^\sharp(d) \vee \dots \vee \llbracket s_k \rrbracket^\sharp(d)$, Lemma 5.1, can be applied to provide us with the desired result. \square

The problem for arbitrary statement is more difficult. By clear equivalence with satisfiability solving modulo the theory of linear real arithmetic, we obtain:

Lemma 5.3. *The problem of deciding, whether or not, for a given template constraint matrix T , and a given statement s , $\llbracket s \rrbracket^\sharp(\infty) > \underline{-\infty}$ holds, is **NP**-complete. \square*

5.2. A Trivial Method for Computing Abstract Semantics. Using the results we have obtained so far, the abstract semantics of a program G w.r.t. some template constraint matrix T can be computed using the following two-step procedure:

- (1) Replace each statement s of the program G with an equivalent merge-simple statement. This corresponds to an explicit enumeration of all paths between cut-points, which potentially causes an exponential blowup.
- (2) Apply the methods of Gawlitza and Seidl [26] to the obtained program to compute the abstract semantics V^\sharp of G .

Because of the possible exponential blowup, the above described method is impractical for most cases⁵. Our method eschews this blowup as follows: instead of enumerating all program paths, we shall visit them only as needed. Guided by a SAT modulo linear real arithmetic solver, our method selects a path through a statement s only when it is *locally profitable* in some sense. In the worst case, an exponential number of paths may be visited (Section 7.3); but one can hope that this rarely happens in practice. In cases in which our algorithm needs exponential time, it at least avoids the explicit exponential expansions. It uses only polynomial space.

⁵ Note that we cannot expect a polynomial-time algorithm, because of Lemma 5.3: even without loops, abstract reachability is **NP**-hard. Even if all statements are merge-simple, we cannot expect a polynomial-time algorithm, since the problem of computing the winning regions of parity games is polynomial-time reducible to abstract reachability [27].

6. MAX-STRATEGY ITERATION

This section presents our main contribution. We adapt the max-strategy improvement schema of Gawlitza and Seidl [28] to obtain an algorithm to compute abstract semantics in the framework of this article.

6.1. Notations. Before we go in medias res, we have to introduce some notations. A system \mathcal{E} of (fixpoint) equations over $\overline{\mathbb{R}}$ is a finite set $\{\mathbf{x}_1 = e_1, \dots, \mathbf{x}_n = e_n\}$ of equations. Here, $\mathbf{x}_1, \dots, \mathbf{x}_n$ are pairwise distinct, $\overline{\mathbb{R}}$ -valued variables and e_1, \dots, e_n are expressions over $\overline{\mathbb{R}}$. We denote the set $\{\mathbf{x}_1, \dots, \mathbf{x}_n\}$ of variables of \mathcal{E} by $\mathbf{X}_{\mathcal{E}}$. We omit the subscript, whenever it is clear from the context. A function $\rho : \mathbf{X} \rightarrow \overline{\mathbb{R}}$ is called a *variable assignment*. It assigns the value $\rho(\mathbf{x})$ to each variable $\mathbf{x} \in \mathbf{X}$. Variable assignments are ordered by the point-wise extension of \leq on $\overline{\mathbb{R}}$, i.e., $\rho \leq \rho'$ if and only if $\rho(\mathbf{x}) \leq \rho'(\mathbf{x})$ for all $\mathbf{x} \in \mathbf{X}$. Since $\overline{\mathbb{R}}$ is a complete linearly ordered set, the set $\mathbf{X} \rightarrow \overline{\mathbb{R}}$ of all variable assignments is a complete lattice. The semantics $\llbracket e \rrbracket : (\mathbf{X} \rightarrow \overline{\mathbb{R}}) \rightarrow \overline{\mathbb{R}}$ of an expression e is defined by $\llbracket \mathbf{x} \rrbracket(\rho) := \rho(\mathbf{x})$ and $\llbracket f(e_1, \dots, e_k) \rrbracket(\rho) := f(\llbracket e_1 \rrbracket(\rho), \dots, \llbracket e_k \rrbracket(\rho))$, where $\mathbf{x} \in \mathbf{X}$, f is a k -ary operator on $\overline{\mathbb{R}}$, e_1, \dots, e_k are expressions, and $\rho : \mathbf{X} \rightarrow \overline{\mathbb{R}}$ is a variable assignment. We define the operator $\llbracket \mathcal{E} \rrbracket : (\mathbf{X} \rightarrow \overline{\mathbb{R}}) \rightarrow \mathbf{X} \rightarrow \overline{\mathbb{R}}$ by $\llbracket \mathcal{E} \rrbracket(\rho)(\mathbf{x}) := \llbracket e \rrbracket \rho$ for all equations $\mathbf{x} = e$ of \mathcal{E} , all $\rho : \mathbf{X} \rightarrow \overline{\mathbb{R}}$, and all $\mathbf{x} \in \mathbf{X}$. A fixpoint equation $\mathbf{x} = e$ is called *monotone* if and only if all operators used in e are monotone. Then, the evaluation function $\llbracket e \rrbracket$ of e is monotone, too. Finally, the operator $\llbracket \mathcal{E} \rrbracket$ is monotone for all systems \mathcal{E} of monotone (fixpoint) equations. A variable assignment ρ is called a *solution* (resp. *pre-solution*, resp. *post-solution*) of \mathcal{E} if and only if $\rho = \llbracket \mathcal{E} \rrbracket(\rho)$ (resp. $\rho \leq \llbracket \mathcal{E} \rrbracket(\rho)$, resp. $\rho \geq \llbracket \mathcal{E} \rrbracket(\rho)$). The least solution of \mathcal{E} is denoted by $\mu\llbracket \mathcal{E} \rrbracket$. If the operator $\llbracket \mathcal{E} \rrbracket$ is monotone, then the fixpoint theorem of Knaster/Tarski (Theorem 2.1) ensures the existence of a uniquely determined least solution $\mu\llbracket \mathcal{E} \rrbracket$. For a system \mathcal{E} of equations and a pre-solution ρ , $\mu_{\geq \rho}\llbracket \mathcal{E} \rrbracket$ denotes the least solution of \mathcal{E} among the solutions of \mathcal{E} that are greater than or equal to ρ , i.e., $\mu_{\geq \rho}\llbracket \mathcal{E} \rrbracket = \min\{\rho' \mid \rho' = \llbracket \mathcal{E} \rrbracket(\rho')$ and $\rho' \geq \rho\}$. Again, if the operator $\llbracket \mathcal{E} \rrbracket$ is monotone, then the fixpoint theorem of Knaster/Tarski ensures the existence of $\mu_{\geq \rho}\llbracket \mathcal{E} \rrbracket$, since the set $\{\rho' \mid \rho' \geq \rho\}$ is a complete lattice.

6.2. Rewriting the Abstract Semantic Equations. The first step of our method consists of rewriting our static analysis problem into a system of monotone fixpoint equations over $\overline{\mathbb{R}}$. Assume that $G = (N, E, \mathbf{st})$ is a program that has n variables, and $T \in \mathbb{R}^{m \times n}$ is a template constraint matrix. Recall that (w.r.t. T) the abstract semantics of G is the least solution of the following constraint system (cf. (3.7) in Subsection 3.3):

$$\mathbf{V}^{\#}[\mathbf{st}] \geq \alpha(\mathbb{R}^n) \quad \mathbf{V}^{\#}[v] \geq \llbracket s \rrbracket^{\#}(\mathbf{V}^{\#}[u]) \quad \text{for all } (u, s, v) \in E \quad (6.1)$$


The constraint system has exactly one $\overline{\mathbb{R}}^m$ -valued variable $\mathbf{V}^{\#}[v]$ for each program point $v \in N$. For each program point $v \in N$, we decompose the $\overline{\mathbb{R}}^m$ -valued variable $\mathbf{V}^{\#}[v]$ into m $\overline{\mathbb{R}}$ -valued variables $\mathbf{d}_{v,1}, \dots, \mathbf{d}_{v,m}$. That is, we set $(\mathbf{d}_{v,1}, \dots, \mathbf{d}_{v,m})^{\top} = \mathbf{V}^{\#}[v]$. We obtain the following constraint system:

$$\mathbf{d}_{\mathbf{st},i} \geq \infty \quad \text{for all } i \in \{1, \dots, m\} \quad (6.2)$$

$$\mathbf{d}_{v,i} \geq \llbracket s \rrbracket_i^{\#}(\mathbf{d}_{u,1}, \dots, \mathbf{d}_{u,m}) \quad \text{for all } (u, s, v) \in E \text{ and all } i \in \{1, \dots, m\} \quad (6.3)$$

$$\begin{aligned}
G &= (N, E, \mathbf{st}) \\
N &= \{\mathbf{st}, 1\} \\
E &= \{(\mathbf{st}, \mathbf{x}'_1 = 0, 1), (1, s, 1)\} \\
s &\equiv \Phi \wedge (\Phi_1 \vee \Phi_2) \\
\Phi &\equiv \mathbf{x}_1 \leq 1000 \wedge \mathbf{x}'_2 = -\mathbf{x}_1 \\
\Phi_1 &\equiv \mathbf{x}'_2 \leq -1 \wedge \mathbf{x}'_1 = -2\mathbf{x}_1 \\
\Phi_2 &\equiv -\mathbf{x}'_2 \leq 0 \wedge \mathbf{x}'_1 = -\mathbf{x}_1 + 1
\end{aligned}$$

(a) The program G



$$T = \begin{pmatrix} 1 & 0 \\ -1 & 0 \end{pmatrix}$$

(b) The template constraint matrix T

(only x_1 is taken into account in the template, thus the zero right column)

$$\begin{aligned}
\mathbf{d}_{\mathbf{st},1} &= \infty & \mathbf{d}_{1,1} &= \max \left\{ \llbracket \mathbf{x}'_1 = 0 \rrbracket_1^\#(\mathbf{d}_{\mathbf{st},1}, \mathbf{d}_{\mathbf{st},2}), \llbracket s \rrbracket_1^\#(\mathbf{d}_{1,1}, \mathbf{d}_{1,2}) \right\} \\
\mathbf{d}_{\mathbf{st},2} &= \infty & \mathbf{d}_{1,2} &= \max \left\{ \llbracket \mathbf{x}'_1 = 0 \rrbracket_2^\#(\mathbf{d}_{\mathbf{st},1}, \mathbf{d}_{\mathbf{st},2}), \llbracket s \rrbracket_2^\#(\mathbf{d}_{1,1}, \mathbf{d}_{1,2}) \right\}
\end{aligned}$$

(c) The equation system $\mathcal{E}(G, T)$

Figure 4: The running example

The fixpoint theorem of Knaster/Tarski (Theorem 2.1) ensures that the least solution of the above system of inequalities is the least solution of the following equation system:

$$\mathbf{d}_{\mathbf{st},i} = \infty \quad \text{for all } i \in \{1, \dots, m\} \quad (6.4)$$

$$\mathbf{d}_{v,i} = \max \left\{ \llbracket s \rrbracket_i^\#(\mathbf{d}_{u,1}, \dots, \mathbf{d}_{u,m}) \mid (u, s, v) \in E \right\} \quad \text{for all } v \in N \setminus \{\mathbf{st}\}, i \in \{1, \dots, m\} \quad (6.5)$$

We denote the above system of fixpoint equations by $\mathcal{E}(G, T)$. From Section 5, we know that the right-hand sides of $\mathcal{E}(G, T)$ are point-wise maxima of finitely many point-wise minima of finitely many weak-affine operators. We summarize the properties of $\mathcal{E}(G, T)$:

Lemma 6.1. *Let G be a program and $V^\#$ its abstract semantics (w.r.t. the template constraint matrix $T \in \mathbb{R}^{m \times n}$). Let $\rho^* := \mu[\mathcal{E}(G, T)]$ be the least solution of $\mathcal{E}(G, T)$. Then $V_i^\#[v] = \rho^*(\mathbf{d}_{v,i})$ for all program points $v \in N$ and all $i \in \{1, \dots, m\}$. The right-hand sides of $\mathcal{E}(G, T)$ are point-wise maxima of finitely many point-wise minima of finitely many weak-affine operators. Thus, they are in particular point-wise maxima of finitely many monotone and concave functions. \square*

Examples 6.2. We again consider our running example specified in Figure 4(a). We want to perform the analysis w.r.t. the template constraint matrix T specified in Figure 4(b). The resulting equation system $\mathcal{E}(G, T)$ is shown in Figure 4(c).

The least solution $\rho^* := \mu[\mathcal{E}(G, T)]$ of $\mathcal{E}(G, T)$ is given by $\rho^* = \{\mathbf{d}_{\mathbf{st},1} \mapsto \infty, \mathbf{d}_{\mathbf{st},2} \mapsto \infty, \mathbf{d}_{1,1} \mapsto 2001, \mathbf{d}_{1,2} \mapsto 2000\}$. Thus, by Lemma 6.1, $V^\#[\mathbf{st}] = (\infty, \infty)$, and $V^\#[1] = (2001, 2000)$. In consequence, all possible values of the program variable \mathbf{x}_1 at program point 1 are in the interval $[-2000, 2001]$. \square

6.3. Adapting the Max-Strategy Improvement Algorithm. Following the lines of Gawlitza and Seidl [29], our starting point is a system \mathcal{E} of monotone fixpoint equations of the form $\mathbf{x} = \max \Sigma_{\mathbf{x}}$, where \mathbf{x} is a $\overline{\mathbb{R}}$ -valued variable, and $\Sigma_{\mathbf{x}}$ is a finite set of monotone and concave expressions over $\overline{\mathbb{R}}$. An expression e is called monotone (resp. concave) if and only if $\llbracket e \rrbracket$ is monotone (resp. concave).⁶ We treat a function from the finite set \mathbf{X} of variables to $\overline{\mathbb{R}}$ as a vector of $|\mathbf{X}|$ elements from $\overline{\mathbb{R}}$. In our application — recall that we aim at solving the equation system $\mathcal{E}(G, T)$ — the sets $\Sigma_{\mathbf{x}}$ are implicitly and succinctly given by the right-hand sides of equations of the forms (6.4) and (6.5). Indeed, every expression of the form $\llbracket s \rrbracket_i^\sharp(\mathbf{d}_{u,1}, \dots, \mathbf{d}_{u,m})$, found on the right-hand side of such equations, can be equivalently rewritten into $\max \{ \llbracket s_1 \rrbracket_i^\sharp(\mathbf{d}_{u,1}, \dots, \mathbf{d}_{u,m}), \dots, \llbracket s_k \rrbracket_i^\sharp(\mathbf{d}_{u,1}, \dots, \mathbf{d}_{u,m}) \}$, where s_1, \dots, s_k are (potentially exponentially many) sequential statements. Since s_1, \dots, s_k are sequential, the operators $\llbracket s_1 \rrbracket_i^\sharp, \dots, \llbracket s_k \rrbracket_i^\sharp$ are point-wise minima of finitely many monotone and weak-affine operators; hence they are monotone and concave operators.

One obvious way to solve the system \mathcal{E} of equations is to perform the above mentioned rewriting explicitly and then apply the max-strategy improvement algorithm. To avoid this impractical exponential blowup, in what follows we modify the algorithm such that it directly works on the succinct representation.

Assume that \mathcal{E} denotes a system of fixpoint equations of the form $\mathbf{x} = \max \Sigma_{\mathbf{x}}$, where $\Sigma_{\mathbf{x}}$ is a finite set of monotone and concave expressions over $\overline{\mathbb{R}}$. A *max-strategy* σ for \mathcal{E} is a system of equations such that, for each equation $\mathbf{x} = e$ of σ , one of the following statements holds:

- (1) e is $-\infty$.
- (2) $e \in \Sigma_{\mathbf{x}}$, where $\mathbf{x} = \max \Sigma_{\mathbf{x}}$ is an equation of \mathcal{E} .

Intuitively, a max-strategy picks for each maximum operator one of its operands. For a system \mathcal{E} of equations, we denote the set of all max-strategies by $\Sigma_{\mathcal{E}}$. In our application, the cardinality of $\Sigma_{\mathcal{E}}$ is exponential in the size of \mathcal{E} . To be more precise, it is in $\mathcal{O}(2^{n^2})$, where n denotes the size of \mathcal{E} . Enumerating all max-strategies is therefore impractical.

Examples 6.3. We continue our running example (Figure 4). Consider the system $\mathcal{E}(G, T)$ and note that $s \equiv \Phi \wedge (\Phi_1 \vee \Phi_2) \equiv (\Phi \wedge \Phi_1) \vee (\Phi \wedge \Phi_2)$; therefore the equation

$$\mathbf{d}_{1,1} = \max \left\{ \llbracket \mathbf{x}'_1 = 0 \rrbracket_1^\sharp(\mathbf{d}_{\text{st},1}, \mathbf{d}_{\text{st},2}), \llbracket s \rrbracket_1^\sharp(\mathbf{d}_{1,1}, \mathbf{d}_{1,2}) \right\} \quad (6.6)$$

can be equivalently rewritten into

$$\mathbf{d}_{1,1} = \max \left\{ \llbracket \mathbf{x}'_1 = 0 \rrbracket_1^\sharp(\mathbf{d}_{\text{st},1}, \mathbf{d}_{\text{st},2}), \llbracket \Phi \wedge \Phi_1 \rrbracket_1^\sharp(\mathbf{d}_{1,1}, \mathbf{d}_{1,2}), \llbracket \Phi \wedge \Phi_2 \rrbracket_1^\sharp(\mathbf{d}_{1,1}, \mathbf{d}_{1,2}) \right\}. \quad (6.7)$$

Recall that this expansion is solely for the purpose of proving properties: it is not done in the algorithm. The equation system σ consisting of the equations

$$\mathbf{d}_{\text{st},1} = \infty \quad \mathbf{d}_{1,1} = \llbracket \Phi \wedge \Phi_2 \rrbracket_1^\sharp(\mathbf{d}_{1,1}, \mathbf{d}_{1,2}) \quad \mathbf{d}_{\text{st},2} = \infty \quad \mathbf{d}_{1,2} = \llbracket \mathbf{x}'_1 = 0 \rrbracket_2^\sharp(\mathbf{d}_{\text{st},1}, \mathbf{d}_{\text{st},2}) \quad (6.8)$$

is thus a max-strategy for this system. \square

⁶For a precise definition of concavity for functions from the set $\overline{\mathbb{R}}^n \rightarrow \overline{\mathbb{R}}^m$, we refer to Gawlitza and Seidl [31]. For this article, however, a precise treatment of these issues is not required. We just mention concavity to give a better intuition.

A crucial notion we need in the following is the notion of improvements. Let σ be a max-strategy for \mathcal{E} and ρ a pre-solution of σ . A max-strategy σ' for \mathcal{E} is called an *improvement* of σ w.r.t. ρ if and only if the following conditions are fulfilled:

- (1) If $\rho \neq \llbracket \mathcal{E} \rrbracket(\rho)$, then $\llbracket \sigma' \rrbracket(\rho) > \rho$.
- (2) If $\mathbf{x} = e$ is an equation of σ and $\mathbf{x} = e'$ is an equation of σ' with $e \neq e'$, then $\llbracket e' \rrbracket(\rho) > \llbracket e \rrbracket(\rho)$.

Example 6.4. We continue our running example (Figure 4). We consider the equation system σ' that consists of the following equations:

$$\mathbf{d}_{\text{st},1} = \infty \quad \mathbf{d}_{1,1} = \llbracket \Phi \wedge \Phi_2 \rrbracket_1^\sharp(\mathbf{d}_{1,1}, \mathbf{d}_{1,2}) \quad \mathbf{d}_{\text{st},2} = \infty \quad \mathbf{d}_{1,2} = \llbracket \Phi \wedge \Phi_1 \rrbracket_2^\sharp(\mathbf{d}_{1,1}, \mathbf{d}_{1,2}) \quad (6.9)$$

The equation system σ' is a max-strategy of $\mathcal{E}(G, T)$ and moreover an improvement of the max-strategy σ (defined in Example 6.3) w.r.t. the variable assignment

$$\rho := \{\mathbf{d}_{\text{st},1} \mapsto \infty, \mathbf{d}_{\text{st},2} \mapsto \infty, \mathbf{d}_{1,1} \mapsto 1, \mathbf{d}_{1,2} \mapsto 0\}. \quad (6.10)$$

It is an improvement, since $\llbracket \llbracket \Phi \wedge \Phi_1 \rrbracket_2^\sharp(\mathbf{d}_{1,1}, \mathbf{d}_{1,2}) \rrbracket(\rho) = 1 > 0 = \llbracket \llbracket \mathbf{x}'_1 = 0 \rrbracket_2^\sharp(\mathbf{d}_{1,1}, \mathbf{d}_{1,2}) \rrbracket(\rho)$. In this example, σ' is the only improvement of σ w.r.t. ρ . \square

Note that, for a max-strategy σ and a pre-solution ρ of σ , there might be several max-strategies σ' that are improvements of σ w.r.t. ρ . Consider, for instance, the equation system $\mathcal{E} = \{\mathbf{x} = \max\{0, 1, 2\}\}$. Both, the max-strategies $\{\mathbf{x} = 1\}$ and $\{\mathbf{x} = 2\}$ are improvements of the max-strategy $\{\mathbf{x} = 0\}$. For the results we are going to develop in this article, it is not important which improvement we choose: this will neither affect the final result obtained, nor change the worst-case complexity bounds that we prove. It is however possible that different heuristics may lead to different practical complexities.

The max-strategy improvement algorithm starts with the max-strategy $\sigma_0 := \{\mathbf{x} = -\infty \mid \mathbf{x} \in \mathbf{X}\}$ and the variable assignment $\rho_0 := \{\mathbf{x} \mapsto -\infty \mid \mathbf{x} \in \mathbf{X}\}$. The algorithm successively performs the following two steps in the given order until it has found the least solution:

- (1) Improve the max-strategy σ w.r.t. ρ .
- (2) Evaluate the max-strategy σ w.r.t. ρ to obtain a new value for ρ .

In pseudo-code, we can formulate it as follows:

Algorithm 1 The Max-Strategy Improvement Algorithm

```

1 :  $\sigma \leftarrow \sigma_0$ ;
2 :  $\rho \leftarrow \rho_0$ ;
3 : while ( $\rho < \llbracket \mathcal{E} \rrbracket(\rho)$ ) {
4 :    $\sigma \leftarrow$  improvement of  $\sigma$  w.r.t.  $\rho$ ;
5 :    $\rho \leftarrow \mu_{\geq \rho} \llbracket \sigma \rrbracket$ ;
6 : }
7 : return  $\rho$ ;
```

For all $i \in \mathbb{N}$, let ρ_i be the value of the variable ρ and σ_i be the value of the variable σ after the i -th evaluation of the loop-body. We have:

Lemma 6.5 ([31], [28, Lem. 6.7]). *The following statements hold for all $i \in \mathbb{N}$:*

- (1) $\rho_i \leq \mu \llbracket \mathcal{E} \rrbracket$.

(2) $\rho_i \leq \llbracket \sigma_{i+1} \rrbracket(\rho_i)$.

(3) If $\rho_i < \mu \llbracket \mathcal{E} \rrbracket$, then $\rho_{i+1} > \rho_i$.

(4) If $\rho_i = \mu \llbracket \mathcal{E} \rrbracket$, then $\rho_{i+1} = \rho_i$. □

The above lemma implies that the algorithm returns the least solution, whenever it terminates. Whether or not it terminates depends on the properties of the class of fixpoint equation systems under consideration. In our application, we aim at computing the least solution of the equation system $\mathcal{E}(G, T)$ (see Subsection 6.2). By Lemma 6.1, the right-hand sides of $\mathcal{E}(G, T)$ are point-wise maxima of finitely many monotone and concave functions. More specifically, the right-hand sides are point-wise maxima of finitely many point-wise minima of finitely many weak-affine operators. This property guaranties the termination of the max-strategy improvement algorithm [31][28, §6.1]. At the latest, it terminates after considering each max-strategy at most linearly often (see Lemma 6.8). Before we explain the remaining building blocks, i.e., how to execute program lines 4 and 5, we consider an example.

Example 6.6. We consider our running example. That is, we aim at computing the least solution of the equation system $\mathcal{E}(G, T)$ shown in Figure 4. Running the algorithm can, for instance, give us the following trace:

$$\sigma_0 := \{\mathbf{d}_{\text{st},1} = -\infty, \mathbf{d}_{\text{st},2} = -\infty, \mathbf{d}_{1,1} = -\infty, \mathbf{d}_{1,2} = -\infty\} \quad (6.11)$$

$$\rho_0 := \{\mathbf{d}_{\text{st},1} \mapsto -\infty, \mathbf{d}_{\text{st},2} \mapsto -\infty, \mathbf{d}_{1,1} \mapsto -\infty, \mathbf{d}_{1,2} \mapsto -\infty\} \quad (6.12)$$

$$\sigma_1 := \{\mathbf{d}_{\text{st},1} = \infty, \mathbf{d}_{\text{st},2} = \infty, \mathbf{d}_{1,1} = -\infty, \mathbf{d}_{1,2} = -\infty\} \quad (6.13)$$

$$\rho_1 := \{\mathbf{d}_{\text{st},1} \mapsto \infty, \mathbf{d}_{\text{st},2} \mapsto \infty, \mathbf{d}_{1,1} \mapsto -\infty, \mathbf{d}_{1,2} \mapsto -\infty\} \quad (6.14)$$

$$\sigma_2 := \{\mathbf{d}_{\text{st},1} = \infty, \mathbf{d}_{\text{st},2} = \infty, \mathbf{d}_{1,1} = \llbracket \mathbf{x}'_1 = 0 \rrbracket_1^\sharp(\mathbf{d}_{\text{st},1}, \mathbf{d}_{\text{st},2}), \quad (6.15)$$

$$\mathbf{d}_{1,2} = \llbracket \mathbf{x}'_1 = 0 \rrbracket_2^\sharp(\mathbf{d}_{\text{st},1}, \mathbf{d}_{\text{st},2})\} \quad (6.16)$$

$$\rho_2 := \{\mathbf{d}_{\text{st},1} \mapsto \infty, \mathbf{d}_{\text{st},2} \mapsto \infty, \mathbf{d}_{1,1} \mapsto 0, \mathbf{d}_{1,2} \mapsto 0\} \quad (6.17)$$

$$\sigma_3 := \{\mathbf{d}_{\text{st},1} = \infty, \mathbf{d}_{\text{st},2} = \infty, \mathbf{d}_{1,1} = \llbracket \Phi \wedge \Phi_2 \rrbracket_1^\sharp(\mathbf{d}_{1,1}, \mathbf{d}_{1,2}), \quad (6.18)$$

$$\mathbf{d}_{1,2} = \llbracket \mathbf{x}'_1 = 0 \rrbracket_2^\sharp(\mathbf{d}_{\text{st},1}, \mathbf{d}_{\text{st},2})\} \quad (6.19)$$

$$\rho_3 := \{\mathbf{d}_{\text{st},1} \mapsto \infty, \mathbf{d}_{\text{st},2} \mapsto \infty, \mathbf{d}_{1,1} \mapsto 1, \mathbf{d}_{1,2} \mapsto 0\} \quad (6.20)$$

$$\sigma_4 := \{\mathbf{d}_{\text{st},1} = \infty, \mathbf{d}_{\text{st},2} = \infty, \mathbf{d}_{1,1} = \llbracket \Phi \wedge \Phi_2 \rrbracket_1^\sharp(\mathbf{d}_{1,1}, \mathbf{d}_{1,2}), \quad (6.21)$$

$$\mathbf{d}_{1,2} = \llbracket \Phi \wedge \Phi_1 \rrbracket_2^\sharp(\mathbf{d}_{1,1}, \mathbf{d}_{1,2})\} \quad (6.22)$$

$$\rho_4 := \{\mathbf{d}_{\text{st},1} \mapsto \infty, \mathbf{d}_{\text{st},2} \mapsto \infty, \mathbf{d}_{1,1} \mapsto 2001, \mathbf{d}_{1,2} \mapsto 2000\} \quad (6.23)$$

Here, for all i , $\rho_{i+1} = \mu_{\geq \rho_i} \llbracket \sigma_{i+1} \rrbracket$ and σ_{i+1} is an improvement of σ_i w.r.t. ρ_i . The variable ρ_4 is a solution of $\mathcal{E}(G, T)$. The max-strategy improvement algorithm terminates with the correct least solution, which is ρ_4 . □

We now present methods to evaluate max-strategies (Line 5 of Algorithm 1) and to improve max-strategies (Line 4 of Algorithm 1).

6.4. Evaluating Max-Strategies. We restrict our consideration to our application. That is, we assume that the equation system \mathcal{E} is given by $\mathcal{E} = \mathcal{E}(G, T)$ for some program G and some template constraint matrix T . For all $i \in \mathbb{N}$, this allows us to compute ρ_i as follows:

Lemma 6.7 ([31],[28]). *Let $i \in \mathbb{N}$. Recall that, by construction, $\rho_{i+1} = \mu_{\geq \rho_i} \llbracket \sigma_{i+1} \rrbracket$. The variable assignment ρ_{i+1} can be computed as follows: Let \mathcal{E}' denote the system of equations that is obtained from the equation system σ_{i+1} by performing the following steps:*

- (1) *Remove every equation $\mathbf{x} = e$, where $\llbracket e \rrbracket(\rho_i) = -\infty$ and replace then the remaining occurrences of \mathbf{x} with the constant $-\infty$.*
- (2) *Remove every equation $\mathbf{x} = e$, where $\llbracket e \rrbracket(\rho_i) = \infty$ and replace then the remaining occurrences of \mathbf{x} with the constant ∞ .*

For all equations $\mathbf{x} = e$ of the equation system σ_{i+1} with $-\infty < \llbracket e \rrbracket(\rho_i) < \infty$, we can compute $\rho_{i+1}(\mathbf{x})$ as follows:

$$\rho_{i+1}(\mathbf{x}) = \sup \{ \rho(\mathbf{x}) \mid \rho : \mathbf{X}_{\mathcal{E}'} \rightarrow \mathbb{R}, \rho \leq \llbracket \mathcal{E}' \rrbracket(\rho) \} \quad (6.24)$$

The value ρ_{i+1} only depends on the equation system σ_{i+1} and the set of variables already identified to be ∞ , namely, $\{ \mathbf{x} \mid \mathbf{x} = e \text{ is an equation of } \sigma_{i+1} \text{ with } \llbracket e \rrbracket(\rho_i) = \infty \}$. \square

In consequence, the max-strategy improvement algorithm has to consider each max-strategy at most $|\mathbf{X}|$ times. Hence, we have:

Lemma 6.8 ([31],[28]). *The max-strategy improvement algorithm terminates after at most $|\mathbf{X}| \cdot |\Sigma_{\mathcal{E}}|$ max-strategy improvement steps. \square*

Lemma 6.7 gives us a method for computing ρ_i . For each variable $\mathbf{x} \in \mathbf{X}$, we have to compute

$$\sup \{ \rho(\mathbf{x}) \mid \rho : \mathbf{X}_{\mathcal{E}'} \rightarrow \mathbb{R} \text{ and } \rho \leq \llbracket \mathcal{E}' \rrbracket(\rho) \}. \quad (6.25)$$

The equations of \mathcal{E}' are of the form $\mathbf{b} = \llbracket s \rrbracket_k^\sharp(\mathbf{b}_1, \dots, \mathbf{b}_m)$, where $\mathbf{b}, \mathbf{b}_1, \dots, \mathbf{b}_m$ are $\overline{\mathbb{R}}$ -valued variables, and s is a sequential statement. Thus, by Lemma 5.1, the right-hand sides are point-wise minima of finitely many monotone and weak-affine functions. Hence, they are monotone and concave. Therefore, (6.25) represents a convex optimization problem.

The above convex optimization problem is of a very special form. The right-hand sides are parameterized linear programs. In consequence, the convex optimization problem can be rewritten into an equivalent linear programming problem as follows: In accordance to (5.1) and (5.2), in \mathcal{E}' , we replace each equation $\mathbf{b} = \llbracket s \rrbracket_k^\sharp(\mathbf{b}_1, \dots, \mathbf{b}_m)$ with the following linear constraints:

$$\mathbf{b} \leq T_k(\mathbf{y}'_1, \dots, \mathbf{y}'_n)^\top \quad (6.26)$$

$$\Phi \quad (6.27)$$

$$T(\mathbf{y}_1, \dots, \mathbf{y}_n)^\top \leq (\mathbf{b}_1, \dots, \mathbf{b}_m) \quad (6.28)$$

Here, $\mathbf{y}_1, \dots, \mathbf{y}_n, \mathbf{y}'_1, \dots, \mathbf{y}'_n$ are fresh variables. Φ is a set of linear inequalities that is obtained from the sequential statement s by

- (1) replacing the variables $\mathbf{x}_1, \dots, \mathbf{x}_n, \mathbf{x}'_1, \dots, \mathbf{x}'_n$ with the fresh variables $\mathbf{y}_1, \dots, \mathbf{y}_n, \mathbf{y}'_1, \dots, \mathbf{y}'_n$,
- (2) replacing all other variables of s with fresh variables, and
- (3) replacing every strict inequality $<$ with a non-strict inequality \leq .

We denote the resulting constraint system by \mathcal{C} . By construction, we have:

$$\sup \{ \rho(\mathbf{x}) \mid \rho : \mathbf{X} \rightarrow \mathbb{R} \text{ and } \rho \leq \llbracket \mathcal{E}' \rrbracket(\rho) \} = \sup \{ \rho(\mathbf{x}) \mid \rho : \mathbf{X} \rightarrow \mathbb{R} \text{ and } \rho \text{ solves } \mathcal{C} \} \quad (6.29)$$

The construction can be carried out in polynomial time. Since \mathcal{C} is a set of linear constraints, we can use linear programming to compute the optimal value. We have:

Lemma 6.9 (Evaluating Max-Strategies). *Whenever our max-strategy improvement algorithm has to compute $\mu_{\geq \rho}[\sigma]$, this can be performed by solving $|\mathbf{X}|$ linear programming problems of polynomial size. The linear programming problems do only depend on σ and the set $\{\mathbf{x} \mid \mathbf{x} = e \text{ is an equation of } \sigma \text{ with } \llbracket e \rrbracket(\rho) = \infty\}$. \square*

Example 6.10. We now discuss how to compute $\rho_3 := \mu_{\geq \rho_2}[\sigma_3]$ from Example 6.6. Note that the values of the variables $\mathbf{d}_{\text{st},1}$ and $\mathbf{d}_{\text{st},2}$ are already known to be ∞ . It remains to determine the values for the variables $\mathbf{d}_{1,1}$ and $\mathbf{d}_{1,2}$. According to Lemma 6.7, we have

$$\begin{aligned} \rho_3(\mathbf{d}_{1,1}) = \sup \{ & \mathbf{d}_{1,1} \mid \mathbf{d}_{1,1}, \mathbf{d}_{1,2} \in \mathbb{R}, \mathbf{d}_{1,1} \leq \llbracket \Phi \wedge \Phi_2 \rrbracket_1^\sharp(\mathbf{d}_{1,1}, \mathbf{d}_{1,2}), \\ & \mathbf{d}_{1,2} \leq \llbracket \mathbf{x}'_1 = 0 \rrbracket_2^\sharp(\infty, \infty) \} \end{aligned} \quad (6.30)$$

Observe that $\Phi \wedge \Phi_2$ can be equivalently rewritten into $\mathbf{x}_1 \leq 0 \wedge \mathbf{x}'_1 = -\mathbf{x}_1 + 1 \wedge \mathbf{x}'_2 = -\mathbf{x}_1$. Thus, according to the above observations, $\rho_3(\mathbf{d}_{1,1})$ is the optimal value of the following linear programming problem:

$$\max \mathbf{d}_{1,1} \quad \mathbf{d}_{1,1} \leq -\mathbf{x}_1 + 1 \quad \mathbf{x}_1 \leq 0 \quad \mathbf{x}_1 \leq \mathbf{d}_{1,1} \quad -\mathbf{x}_1 \leq \mathbf{d}_{1,2} \quad \mathbf{d}_{1,2} \leq 0 \quad (6.31)$$

Since the optimal value is 1, we get $\rho_3(\mathbf{d}_{1,1}) = 1$. Similarly, to compute $\rho_3(\mathbf{d}_{1,2})$, we compute the optimal value of the following linear programming problem:

$$\max \mathbf{d}_{1,2} \quad \mathbf{d}_{1,1} \leq -\mathbf{x}_1 + 1 \quad \mathbf{x}_1 \leq 0 \quad \mathbf{x}_1 \leq \mathbf{d}_{1,1} \quad -\mathbf{x}_1 \leq \mathbf{d}_{1,2} \quad \mathbf{d}_{1,2} \leq 0 \quad (6.32)$$

This gives us $\rho_3(\mathbf{d}_{1,2}) = 0$.

Both linear programming problems have the same feasible space. This can be utilized in an implementation to improve the performance. Furthermore, $\rho_3(\mathbf{d}_{1,1}) = \mathbf{d}_{1,1}^*$ and $\rho_3(\mathbf{d}_{1,2}) = \mathbf{d}_{1,2}^*$ for any optimal solution $(\mathbf{d}_{1,1}^*, \mathbf{d}_{1,2}^*, \mathbf{y}_1^*)$ of the following linear programming problem:

$$\max \mathbf{d}_{1,1} + \mathbf{d}_{1,2} \quad \mathbf{d}_{1,1} \leq -\mathbf{x}_1 + 1 \quad \mathbf{x}_1 \leq 0 \quad \mathbf{x}_1 \leq \mathbf{d}_{1,1} \quad -\mathbf{d}_1 \leq \mathbf{d}_{1,2} \quad \mathbf{d}_{1,2} \leq 0 \quad (6.33)$$

Hence, for this example, it is sufficient to solve one linear programming problem to determine the variable assignment ρ_3 . \square

The technique for evaluating max-strategies can thus be further optimized. It is not necessary to solve one linear program for each variable. Instead, it is possible to evaluate a max-strategy entirely by solving only two linear programming problems of linear size. The solution of the first linear programming problem tells us which variables are to set to ∞ . The solution of the second linear programming problem provides us with the values of the variables which receive finite values. In this article, we do not elaborate on these techniques.

6.5. Improving Max-Strategies. We now discuss how we can compute an improvement of a max-strategy σ w.r.t. a variable assignment ρ . Since, by Lemma 5.3, this problem is NP-hard, we cannot expect to come up with a polynomial time algorithm. We propose a solution that utilizes SMT solving techniques.

Let us first explain the intuition of our method, which is very similar to how the “path focusing” technique from Monniaux and Gonnord [49] selects the next iteration path. A strategy needs improvement if and only if its value does not define an inductive invariant. In other words: there is an outgoing transition from the “invariant candidate” into its complement, meaning that there is an execution trace through a statement, starting from the invariant candidate and ending with a violation of the current bounds. Whether this holds is a SAT problem modulo (SMT) the theory of linear real arithmetic; it can therefore

be solved by SMT-solvers. Furthermore, the solution from the SMT problem picks one of the sequential statements from the merge-simple expansion of the statement as “offending”, explaining why the invariant candidate is not an invariant; in other words, it points to a possible improvement in the strategy. More generally, the set of solutions of the SMT problem maps to the possible improvements.

Let us now see this process more formally. Assume that we have to improve a given max-strategy

$$\sigma = \{\mathbf{x}_1 = \sigma_1, \dots, \mathbf{x}_n = \sigma_n\} \quad (6.34)$$

for the equation system

$$\mathcal{E} = \{\mathbf{x}_1 = e_1, \dots, \mathbf{x}_n = e_n\} \quad (6.35)$$

w.r.t. a variable assignment ρ , which is a solution of σ , i.e., $\rho = \llbracket \sigma \rrbracket(\rho)$. This is exactly the situation we are concerned with, when we execute our max-strategy improvement algorithm. For each $i \in \{1, \dots, n\}$, we now want to check whether or not $\rho(\mathbf{x}_i) < \llbracket e_i \rrbracket \rho$. If this is the case, we moreover want to compute a max-strategy σ'_i for e_i such that $\rho(\mathbf{x}_i) < \llbracket \sigma'_i \rrbracket \rho$. Note that, since $\rho(\mathbf{x}_i) < \llbracket e_i \rrbracket \rho$, we could also compute a max-strategy σ'_i such that $\llbracket \sigma'_i \rrbracket \rho = \llbracket e_i \rrbracket \rho$. If $\rho(\mathbf{x}_i) < \llbracket e_i \rrbracket \rho$ does not hold, then we set $\sigma'_i := \sigma_i$. Finally, the max-strategy $\sigma' := \{\mathbf{x}_1 = \sigma'_1, \dots, \mathbf{x}_n = \sigma'_n\}$ is an improvement of σ w.r.t. ρ .

Given an equation $\mathbf{x} = e$ and a variable assignment ρ , we must decide whether or not $\rho(\mathbf{x}) < \llbracket e \rrbracket(\rho)$ holds, and compute a max-strategy σ' of e such that $\rho(\mathbf{x}) < \llbracket \sigma' \rrbracket(\rho)$ holds. Recall that the semantic equations we are concerned with in this article are of the form

$$\mathbf{x} = \max \{e_1, \dots, e_k\} \quad (6.36)$$

where, for all $i \in \{1, \dots, k\}$, each expression e_i is either a constant or an expression of the form $\llbracket s \rrbracket_j^\sharp(\mathbf{x}_1, \dots, \mathbf{x}_m)$. Hence, we can answer the above question by answering the question for each argument e_1, \dots, e_k of the maximum separately. It thus remains to find a method to check whether or not, for a given statement s , a given $j \in \{1, \dots, m\}$, a given $c \in \mathbb{R} \cup \{-\infty\}$, and a given $d \in \overline{\mathbb{R}}^m$, $\llbracket s \rrbracket_j^\sharp(d) > c$ holds — which is, by Lemma 5.3, a **NP**-hard problem. Our approach is to construct the following SAT modulo linear real arithmetic formula (we use existential quantifiers to improve readability):

$$\Psi(s, d, j, c) := \exists \mathbf{v} \in \mathbb{R}. \Psi(s, d, j) \wedge \mathbf{v} > c \quad (6.37)$$

$$\Psi(s, d, j) := \exists \mathbf{x} \in \mathbb{R}^n, \mathbf{x}' \in \mathbb{R}^n. T\mathbf{x} \leq d \wedge \Psi(s) \wedge \mathbf{v} = T_j.\mathbf{x}' \quad (6.38)$$

Here, $\Psi(s)$ is a formula that relates every $x \in \mathbb{R}^n$ with all elements from the set $\llbracket s \rrbracket\{x\}$. It is defined inductively over the structure of the statement s as follows:

$$\Psi(s) := s \quad \text{if } s \text{ is a literal} \quad (6.39)$$

$$\Psi(s_1 \wedge s_2) := \Psi(s_1) \wedge \Psi(s_2) \quad (6.40)$$

$$\Psi(s_1 \vee s_2) := (\neg \mathbf{a}_{s_1 \vee s_2} \wedge \Psi(s_1)) \vee (\mathbf{a}_{s_1 \vee s_2} \wedge \Psi(s_2)) \quad (6.41)$$

Here, for every sub-formula $s_1 \vee s_2$ of s , $\mathbf{a}_{s_1 \vee s_2}$ is a fresh Boolean variable. The set of free variables of the formula $\Psi(s)$ is

$$\{\mathbf{x}, \mathbf{x}'\} \cup \{\mathbf{a}_{s_1 \vee s_2} \mid s_1 \vee s_2 \text{ is a sub-formula of } s\}. \quad (6.42)$$

The variables \mathbf{x} and \mathbf{x}' are \mathbb{R}^n -valued variables. By construction, $s[x/\mathbf{x}, x'/\mathbf{x}']$ is satisfiable if and only if $\Psi(s)[x/\mathbf{x}, x'/\mathbf{x}']$ is satisfiable for all $x, x' \in \mathbb{R}^n$. That is, s and $\Psi(s)$ are describing the same relation. We therefore obtain the following lemma:

Lemma 6.11. $\llbracket s \rrbracket_j^\sharp(d) > c$ if and only if $\Psi(s, d, j, c)$ is satisfiable. \square

The difference between the formula s and the formula $\Psi(s)$ is that the Boolean variables of the formula $\Psi(s)$ additionally describe a path through the formula. More precisely, a valuation for the variables from the set $\{\mathbf{a}_{s_1 \vee s_2} \mid s_1 \vee s_2 \text{ is a sub-formula of } s\}$ describes a path through s .

Let s be a statement, $d \in \overline{\mathbb{R}}^m$, $j \in \{1, \dots, m\}$, and $c \in \mathbb{R} \cup \{-\infty\}$. Assume now that $\llbracket s \rrbracket_j^\sharp(d) > c$. Our next goal is to compute a max-strategy σ for the statement s such that $\llbracket \sigma \rrbracket_j^\sharp(d) > c$. By Lemma 6.11, there exists a model M of $\Psi(s, d, j, c)$. We define the max-strategy σ_M for the statement s recursively by

$$\sigma_M(s) := s \quad \text{if } s \text{ is a literal} \quad (6.43)$$

$$\sigma_M(s_1 \wedge s_2) := \sigma_M(s_1) \wedge \sigma_M(s_2) \quad (6.44)$$

$$\sigma_M(s_1 \vee s_2) := \begin{cases} \sigma_M(s_1) & \text{if } M(\mathbf{a}_{s_1 \vee s_2}) = 0 \\ \sigma_M(s_2) & \text{if } M(\mathbf{a}_{s_1 \vee s_2}) = 1 \end{cases} \quad (6.45)$$

By again applying Lemma 6.11, we get $\llbracket \sigma_M \rrbracket_j^\sharp(d) > c$ and thus the following lemma:

Lemma 6.12. *By solving the SAT modulo linear real arithmetic formula $\Psi(s, d, j, c)$ that can be obtained from s in linear time, we can decide, whether or not $\llbracket s \rrbracket_j^\sharp(d) > c$ holds. From a model M of this formula, we can, in linear time, obtain a \vee -strategy σ_M for s such that $\llbracket \sigma_M \rrbracket_j^\sharp(d) > c$. \square*

Example 6.13. We again continue our running example, which is summarized in Figure 4. We want to know, whether or not $\llbracket s \rrbracket_1^\sharp(0, 0) > 0$ holds. For that we compute a model M of the formula $\Psi(s, (0, 0), 1, 0)$ which is given as follows:

$$\Psi(s, (0, 0), 1, 0) \equiv \exists \mathbf{v} \in \mathbb{R} . \Psi(s, (0, 0)^\top, 1) \wedge \mathbf{v} > 0 \quad (6.46)$$

$$\Psi(s, (0, 0), 1) \equiv \exists \mathbf{x} \in \mathbb{R}^2, \mathbf{x}' \in \mathbb{R}^2 . \mathbf{x}_1 \leq 0 \wedge -\mathbf{x}_1 \leq 0 \wedge \Psi(s) \wedge \mathbf{v} = \mathbf{x}'_1 \quad (6.47)$$

$$\Psi(s) \equiv \Phi \wedge ((-\mathbf{a}_{\Phi_1 \vee \Phi_2} \wedge \Phi_1) \vee (\mathbf{a}_{\Phi_1 \vee \Phi_2} \wedge \Phi_2)) \quad (6.48)$$

The formulas Φ , Φ_1 , and Φ_2 are defined in Figure 4. $M = \{a_{\Phi_1 \vee \Phi_2} \mapsto 1\}$ is a model, which gives us the max-strategy $\sigma_M \equiv \Phi \wedge \Phi_2$ for s . Thus, by Lemma 6.12, we have $\llbracket \sigma_M \rrbracket_1^\sharp(0, 0) = \llbracket \Phi \wedge \Phi_2 \rrbracket_1^\sharp(0, 0) > 0$. \square

We must still provide a method for computing the values for the Boolean variables of a model of the formula $\Psi(s, d, j, c)$. Most of the state-of-the-art SMT solvers, such as Yices [21, 22], support the computation of models directly; the SMTLIB2 standard [6] has a `get-assignment` command that can be used to extract the Boolean part of a model. If this feature is not supported, one can compute the model, or only the values for the Boolean variables, using standard self-reduction techniques.

Recall that the semantic equations we are concerned with in this article are of the form $\mathbf{x} = \max\{e_1, \dots, e_k\}$, where each expression e_i , for all $i \in \{1, \dots, k\}$, is either a constant or an expression of the form $\llbracket s \rrbracket_j^\sharp(\mathbf{x}_1, \dots, \mathbf{x}_m)$ where s is a statement. As discussed above, we can check whether or not $\rho(\mathbf{x}) < \llbracket \max\{e_1, \dots, e_k\} \rrbracket(\rho)$ holds, and if this is the case compute a max-strategy σ' such that $\rho(\mathbf{x}) < \llbracket \sigma' \rrbracket(\rho)$ holds, by solving at most k SAT modulo linear real arithmetic formulas, each of which can be constructed in linear time. Equivalently, instead of running k SMT queries, each obtaining a part of the next strategy,

we can rename Boolean variables of these SMT formulas so that they are distinct and query the conjunction of the resulting formulas.

Lemma 6.14. *Let $\mathbf{x} = e$ be an abstract semantic equation, ρ a variable assignment, and $c \in \overline{\mathbb{R}}$. By solving a single SAT modulo linear real arithmetic formula that can be obtained from e , ρ and c in linear time, we can decide, whether or not $\llbracket e \rrbracket \rho > c$ holds. From a model M of this formula, provided that $\llbracket e \rrbracket \rho > c$ holds, we can in linear time obtain a max-strategy σ_M for e such that $\llbracket \sigma_M \rrbracket \rho > c$. \square*

Remark that we did not discuss how to choose the next max-strategy σ' , except that it should satisfy $\rho(\mathbf{x}) < \llbracket \sigma' \rrbracket(\rho)$ (which is ensured by the SMT-solving step). Indeed, there could be many different suitable σ' s, and the SMT-solver may return any of them. There is however at least one that is *locally optimal*, that is, $\llbracket \sigma' \rrbracket(\rho)$ is maximal, otherwise said $\llbracket \sigma' \rrbracket(\rho) = \llbracket e \rrbracket(\rho)$. Future work should include experiments on the performance impact of using the locally optimal strategies instead of just any strategies.

It is possible to obtain a locally optimal strategy by repeated calls to the SMT-solvers. A naive method would be to query the SMT-solver for a σ'' such that $\llbracket \sigma' \rrbracket(\rho) < \llbracket \sigma'' \rrbracket(\rho)$, then for a σ''' such that $\llbracket \sigma'' \rrbracket(\rho) < \llbracket \sigma''' \rrbracket(\rho)$ and so on until there is no locally better strategy; the last strategy obtained is thus locally optimal. A less naive method would be to take a rough bound $M \geq \llbracket e \rrbracket(\rho)$ and perform binary search in the interval $[\llbracket \sigma' \rrbracket(\rho), M]$: at each step, maintain an interval $[a, b]$ and query whether there exists σ'' such that $\llbracket \sigma' \rrbracket(\rho) \geq \frac{a+b}{2}$; if so, replace a by $\frac{a+b}{2}$ and restart, if not, replace b by $\frac{a+b}{2}$ and restart. The SMT-solving community is now considering the problem of *optimization modulo theory* [58] and we can hope for progress in this respect.

7. COMPLEXITY

In this section, we shall prove that the decision problem associated with our computation is at the second level of the polynomial hierarchy, even if there is a single feedback vertex, a single real variable, and a single constraint in the template. It is therefore unsurprising that our algorithm exhibits exponential complexity in the worst case, by enumerating an exponential number of strategies: we shall then provide an artificial example on which it is the case.

7.1. A Lower Bound on the Complexity. In this section we show that the problem of computing abstract semantics of programs w.r.t. the interval domain is Π_2^P -hard. Π_2^P -hard problems are conjectured to be harder than both NP-complete and coNP-complete problems. For further information regarding the polynomial-time hierarchy see, for instance, Papadimitriou [50], Stockmeyer [61].

Theorem 7.1. *The problem of deciding, whether, for a given program G , a given template constraint matrix T , and a given program point v , $V^\sharp[v] > -\infty$ holds, is Π_2^P -hard.*

The problem remains Π_2^P -hard even if the program variables are abstracted at a single program point and the template constraint matrix T is restricted to a single variable x and a single constraint of the form $x \leq B$.

Proof. We reduce the Π_2^P -complete problem of deciding the truth of a $\forall^* \exists^*$ propositional formula [63] to our static analysis problem. Let

$$\Phi \equiv \forall \mathbf{x}_1, \dots, \mathbf{x}_n. \exists \mathbf{y}_1, \dots, \mathbf{y}_m. \Phi' \tag{7.1}$$

be a formula without free variables, where Φ' is a propositional formula. We consider the analysis of the following pseudo-C program, where n is a constant:

```

x = 0;
while (x < 2n) {
  z = x;
  if (x >= 2n-1) { xn=1; x -= 2n-1; } else { xn=0; }
  :
  if (x >= 21-1) { x1=1; x -= 21-1; } else { x1=0; }
  choose(y1, ..., ym);
  if (Φ'(x1, ..., xn, y1, ..., ym)) {
    x++;
  }
}

```

In intuitive terms: this program initializes the program variable x to 0. Then, it enters a loop: compute into x_1, \dots, x_n the binary decomposition of x , and non-deterministically choose y_1, \dots, y_m . If Φ' is true, it increments x by one and loops, unless x reaches 2^n in which case it terminates; otherwise, it just loops. Thus, there exists a terminating computation if and only if Φ holds.

We reformulate the above pseudo-C program into the program $G = (N, E, \mathbf{st})$ that uses only one program variable \mathbf{x} , where

- (1) $N = \{\mathbf{st}, 1, 2\}$ is the set of program points, and
- (2) $E = \{(\mathbf{st}, \mathbf{x}' = 0, 1), (1, s, 1), (1, \mathbf{x} \geq 2^n, 2)\}$ is the set of control-flow edges, where

$$\begin{aligned}
s &\equiv \mathbf{z}_n = \mathbf{x} \\
&\wedge ((\mathbf{z}_n \geq 2^{n-1} \wedge \mathbf{z}_{n-1} = \mathbf{x} - 2^{n-1} \wedge \mathbf{x}_n = 1) \vee (\mathbf{z}_n \leq 2^{n-1} - 1 \wedge \mathbf{x}_n = 0)) \\
&\wedge \dots \\
&\wedge ((\mathbf{z}_1 \geq 2^{1-1} \wedge \mathbf{z}_0 = \mathbf{x}_1 - 2^{1-1} \wedge \mathbf{x}_1 = 1) \mid (\mathbf{z}_1 \leq 2^{1-1} - 1 \wedge \mathbf{x}_1 = 0)) \\
&\wedge s(\Phi') \\
&\wedge \mathbf{x}' = \mathbf{x} + 1.
\end{aligned}$$

The statement $s(\Phi')$ is obtained by taking formula Φ' in negation normal form (all negations pushed to the leaves), leaving the Boolean structure in place and replacing each positive literal x by a test $x = 1$ and each negative literal $\neg x$ by a test $x = 0$.

With this formalization, Φ holds if and only if $V[2] \neq \emptyset$. For the abstraction, we consider the interval domain, or even simply the domain of upper bounds on \mathbf{x} (i.e., we have constraints of the form $\mathbf{x} \leq b$). By considering the Kleene iteration, it is easy to see that $V[2] \neq \emptyset$ holds if and only if $V^\# [2] > \underline{-\infty}$ holds. Thus Φ holds if and only if $V^\# [2] > \underline{-\infty}$ holds. \square

7.2. An Example with Exponential Running Time Behavior. Recall that the number of strategy improvement steps is exponentially bounded by the size of the input. Each step consists in one phase of SMT-solving for linear real arithmetic followed by solving a linear program of polynomial size. Thus, each step can be performed in exponential time. Therefore, the whole algorithm can be executed in exponential time.

We shall now see that our algorithm takes exponential time on the instances that are similar to the instances generated from the reduction in the proof of Theorem 7.1. The instances generated from the reduction require $\Theta(2^n)$ steps. However, the input is of size $\mathcal{O}(n^2)$, because the numbers $2^{n-1}, 2^{n-2}, \dots, 2^0$ require space $\Theta(n^2)$. We modify the instances generated from the reduction in such a way that the sizes of the programs are in $\mathcal{O}(n)$. We achieve this by introducing auxiliary variables for the numbers $2^{n-1}, 2^{n-2}, \dots, 2^0$. For all $n \in \mathbb{N}$, we define the program $G_n = (N, E, \mathbf{st})$, where

$$N = \{\mathbf{st}, 1\}, \quad (7.2)$$

$$E = \{(\mathbf{st}, \mathbf{x}'_1 = 0, 1), (1, s, 1)\}, \quad (7.3)$$

with

$$s \equiv \mathbf{y}_1 = 1 \wedge \mathbf{y}_2 = 2\mathbf{y}_1 \wedge \dots \wedge \mathbf{y}_n = 2\mathbf{y}_{n-1} \wedge \mathbf{z}_n = \mathbf{x}_1 \quad (7.4)$$

$$\wedge (\mathbf{z}_n \geq \mathbf{y}_n \wedge \mathbf{z}_{n-1} = \mathbf{z}_n - \mathbf{y}_n \vee \mathbf{z}_n \leq \mathbf{y}_n - 1 \wedge \mathbf{z}_{n-1} = \mathbf{z}_n) \quad (7.5)$$

$$\wedge \dots \quad (7.6)$$

$$\wedge (\mathbf{z}_1 \geq \mathbf{y}_1 \wedge \mathbf{z}_0 = \mathbf{z}_1 - \mathbf{y}_1 \vee \mathbf{z}_1 \leq \mathbf{y}_1 - 1 \wedge \mathbf{z}_0 = \mathbf{z}_1) \quad (7.7)$$

$$\wedge \mathbf{x}'_1 = \mathbf{x}_1 + 1. \quad (7.8)$$

Here, \mathbf{x}_1 is the only program variable. It is sufficient to use the template constraint matrix $T = (1)$, which corresponds to the template \mathbf{x}_1 . That is, we are only interested in the upper bound on the value of the variable \mathbf{x}_1 . Remark that the strategy iteration does not depend on the strategy improvement operator in use: at any time there is exactly one possible improvement, until the least solution is reached. All strategies for the statement s will be encountered. Thus, the strategy improvement algorithm performs 2^n strategy improvement steps. Since the size of G_n is $\Theta(n)$, exponentially many strategy improvement steps are performed.

7.3. An Upper Bound on the Complexity. In Section 7.1, we have provided a lower bound on the complexity of computing abstract semantics w.r.t. the template linear domains. The associated decision problem is not only Π_2^p -hard, but in fact Π_2^p -complete:

Theorem 7.2. *The problem of deciding, whether, for a given program G , a given template constraint matrix T , and a given program point v , $V^\sharp[v] > \underline{-\infty}$ holds, is in Π_2^p .*

Proof. We consider the negation of the above problem: for a given program G , a given template constraint matrix T , a given program point v , and a given $i \in \{1, \dots, m\}$, decide whether $V_i^\sharp[v] = -\infty$; we shall now show that this problem is in Σ_2^p .

In non-deterministic polynomial time we can guess a max-strategy σ for $\mathcal{E}' := \mathcal{E}(G, T)$ and a set \mathbf{X}^∞ of variables that have the value ∞ ; these will form the witness for the initial existential quantifier. We can evaluate the max-strategy σ w.r.t. the set of variables \mathbf{X}^∞ assigned to $+\infty$ in polynomial time using linear programming (cf. Subsection 6.4). Let $\rho_{\sigma, \mathbf{X}^\infty}$ denote the resulting variable assignment.

We shall now show that checking whether this strategy (and set of infinite variables) is stable is in **co-NP**. Because of Lemma 5.3, we can use an **NP** oracle to check whether there exists an improvement of the strategy σ w.r.t. $\rho_{\sigma, \mathbf{X}^\infty}$, which is exactly the negation of being stable.

If the strategy is stable, we know that $\rho_{\sigma, \mathbf{X}^\infty} \geq \mu[\mathcal{E}']$ holds. Therefore, by Lemma 6.1, we have $\rho_{\sigma, \mathbf{X}^\infty}(\mathbf{x}_{v,i}) \geq V_i^\sharp[v]$ for all program points $v \in N$ and all $i \in \{1, \dots, m\}$. Since we

n	user time (sec)	number of improvement steps	number of SMT queries	number of linear programs
1	0.10	5	14	8
2	0.17	7	34	12
3	0.38	11	76	20
4	1.02	19	170	36
5	3.64	35	384	68
6	6.97	67	870	132
7	26.02	131	1964	260
8	31.53	259	4402	516
9	95.22	515	9784	1028
10	207.62	1027	21566	2052

Figure 5: Benchmark for the prototypical implementation

also know that there exists some max-strategy σ and some set \mathbf{X}^σ such that $\rho_{\sigma, \mathbf{X}^\infty} = \mu[[\mathcal{E}']]$, we accept, whenever $\rho_{\sigma, \mathbf{X}^\infty}(\mathbf{x}_{v,i}) = -\infty$ holds. \square

8. EXPERIMENTAL RESULTS

We have implemented our presented max-strategy improvement algorithm; our prototype should however be considered as a proof-of-concept. Benchmark results for real examples are left for future work.

The algorithm is implemented in OCaml 3.10.2; it uses Yices 1.0.27 [21, 22] for computing models for SAT modulo linear real arithmetic formulas; for solving the occurring linear programming problems it uses QSOpt-Exact 2.5.6 [3, 23], an exact arithmetic version of QSOpt. We made our experiments under Debian Linux (Lenny) running under Parallels Desktop 4 on an Apple MacBook (2.16 GHz Intel Core 2 Duo, 2GB 667 MHz DDR2 SDRAM). Our solver takes as input a text file that contains the program and the linear templates to be used for the analysis. The benchmark results for the example of Section 7.3 are shown in Figure 5. The number of max-strategy improvement steps grows — as expected — exponentially in n . Briefly, the implementation solves 2 linear programming problems and at most $2(2n+1) = 4n+2$ SMT queries per max-strategy improvement step. The factor 2 comes from the fact that we have 2 program points and the factor $(2n+1)$ from the fact that we have $(2n+1)$ templates. We emphasize that the example is created artificially. Since the problem we are solving is Π_2^P -complete, it is not surprising that there exists an example that does not scale.

For the running example of this article (Example 4.1), our solver computes the correct result after 0.05 seconds.

There are also many possibilities for improving the implementation. On the limited number of examples that we tried with our proof-of-concept implementation, the main computational expense comes from the linear programs that have to be solved. This is mainly due to the fact that we use an exact arithmetic simplex solver and we solve every occurring linear program from scratch although we know beforehand that the linear problems that we have to solve are feasible. Instead of solving each linear program from scratch,

one could use the information obtained from the previously solved linear programs (that are similar). One can also utilize the information obtained from the SMT solver in order to obtain a feasible basis to start the simplex method with.

9. CONCLUSION AND FURTHER RESEARCH DIRECTIONS

We have proposed a method for computing the least fixpoints in template linear constraint domains (e.g., Cartesian products of intervals) of transition systems specified using linear real arithmetic formulas. This allows finding the strongest invariant in this domain of a loop consisting only in linear assignments and non-strict linear inequalities over the real numbers.

Because it distinguishes individual paths in the program, our method does not suffer from the imprecision induced by convex hull operations. These paths are looked up on demand, as results from satisfiability testing, therefore avoiding memory blowup. Our technique, however, has exponential worst case complexity, which is hardly surprising since the decision problem associated with our computation is Π_2^P -complete. Due to limited resources, we have so far not been able to implement it in a tool capable of running on real examples.

It is quite obvious that, due to the use of SMT queries, the size of the problems given as input, and their branching structure, must be limited. One method for limiting the size of the SMT formulas is to decompose the program into statements, thus adding more points at which states are abstracted, as proposed by [49]: this simplifies the problem, but may reduce precision; another method is to restrict the analysis to a subset of the variables, determined by some form of dependency analysis.

The restriction to linear templates and linear statements may seem onerous. It might be possible to apply the same ideas for non-linear templates [30]. With respect to non-linear statements, a possibility is to linearize them [44, 46]: for short, assuming $A \leq x \leq B$ where A and B are constants, then the nonlinear constraint $z = xy$ may be abstracted by the linear constraint $(Ay \leq xy \leq By \wedge y \geq 0) \vee (By \leq xy \leq Ay \wedge y < 0)$. If the assumptions made by the linearization are found not to hold for the fixed point computed by the max-strategy iteration technique, one has to relax these assumptions and restart the solving process.

More generally, one may envision a nesting of two iteration schemes: the inner scheme solving exactly, using max-strategy iteration, a simplification of the concrete program, the outer scheme iterating over possible simplifications. The outer scheme would deal with all program features not supported by our max-strategy iteration algorithm. Consider pointers, for instance: the outer scheme could temporarily assume that x and y may be aliased, while z is not aliased with anything, and then rewrite the program according to these assumptions in order to obtain a pointer-free program (may-alias information becomes non-deterministic choice, while must-aliased variables are merged). This outer iteration may be ascending and optimistic, starting with strong assumptions on the program and relaxing them progressively as the results of the inner scheme invalidate them, or decreasing and pessimistic, starting with weak assumptions and strengthening them progressively as the results of the inner scheme show them to be too severe. Such mixed approaches would cope with programs features not directly supported by our max-strategy iteration solver. Further work is needed in this direction to ascertain which techniques are usable.

Another problem is finding suitable templates — while there exist obvious choices in some cases (intervals for getting rough invariants of control applications, difference bounds

for scheduling applications, etc.), there is no generic method for obtaining good templates. Amato et al. [2] proposed finding templates using principal component analysis, but it is yet unclear whether this approach suited to practical problems. A simple solution may be to run some conventional polyhedral analysis, and keeping the directions of the polyhedra obtained before widening.

Our max-strategy iteration algorithms only deal with real numerical values. We can cope with integers by relaxing them to reals, with the usual precautions ($x < y$ converted to $x \leq y - 1$). Another possible extension is to integrate Boolean types, or more generally finitely enumerated types, into the invariant, or equivalently, to insert them implicitly into the control flow.

An intriguing extension of our framework is the case where the control flow is specified implicitly. The problem considered in this article is expressed as a control-flow graph given by a list of nodes and statements over the transitions. Now consider the addition of n Boolean variables to the system: a common method to encode such variables in a transition system is to distinguish all Boolean combinations and every control node, and thus multiply the number of control nodes by 2^n . Clearly, we would prefer to work directly on the transition relation of the original program, which would include free Boolean variables encoding the departure and arrival control states, and consider our abstract reachability problem on programs expressed using this succinct representation. Since this problem includes Boolean reachability (also known as the reachability problem for succinctly represented graphs), which is **PSPACE**-complete [51], it is **PSPACE**-hard. Our strategy iteration approach can be extended to show that it is in **coNEXPTIME**. We conjecture that it is **coNEXPTIME**-complete, but we have so far not been able to prove it. It is also unknown whether some practically useful algorithms, perhaps based on binary decision diagrams (BDDs), could be devised for this problem.

ACKNOWLEDGMENTS

The authors wish to thank the anonymous referees for their helpful suggestions and references.

REFERENCES

- [1] A. Adjé, S. Gaubert, and E. Goubault. Computing the smallest fixed point of non-expansive mappings arising in game theory and static analysis of programs. preprint, arXiv:0806.1160v3, 2011. <http://arxiv.org/abs/0806.1160>.
- [2] G. Amato, M. Parton, and F. Scozzari. Deriving numerical abstract domains via principal component analysis. In R. Cousot and M. Martel, editors, *Static analysis (SAS)*, volume 6337 of *LNCS*, pages 134–150. Springer, 2010. ISBN 978-3-642-15768-4.
- [3] D. Applegate, W. Cook, S. Dash, and D. Espinoza. Exact solutions to linear programming problems. *Oper. Res. Lett.*, 35(6):693–699, 2007.
- [4] R. Bagnara, P. M. Hill, E. Mazzi, and E. Zaffanella. Widening operators for weakly-relational numeric abstractions. In C. Hankin and I. Siveroni, editors, *Static analysis (SAS)*, volume 3672 of *LNCS*, pages 3–18. Springer, 2005. ISBN 3-540-28584-9. doi: 10.1007/11547662.3.

- [5] R. Bagnara, P. M. Hill, E. Ricci, and E. Zaffanella. Precise widening operators for convex polyhedra. *Science of Computer Programming*, 58(1-2):28–56, 2005. doi: 10.1016/j.scico.2005.02.003.
- [6] C. Barrett, A. Stump, and C. Tinelli. The SMT-LIB Standard: Version 2.0. In A. Gupta and D. Kroening, editors, *SMT*, 2010.
- [7] A. Biere, M. Heule, H. van Maaren, and T. Walsh, editors. *Handbook of satisfiability*, volume 185 of *Frontiers in Artificial Intelligence and Applications*. IOS Press, Amsterdam, 2009. ISBN 978-1-58603-929-5.
- [8] B. Blanchet, P. Cousot, R. Cousot, J. Feret, L. Mauborgne, A. Miné, D. Monniaux, and X. Rival. Design and implementation of a special-purpose static program analyzer for safety-critical real-time embedded software. In *The Essence of Computation: Complexity, Analysis, Transformation*, number 2566 in LNCS, pages 85–108. Springer, 2002.
- [9] 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 *Programming Language Design and Implementation (PLDI)*, pages 196–207. ACM, 2003. ISBN 1-58113-662-5. doi: 10.1145/781131.781153.
- [10] F. Bourdoncle. *Sémantiques des Langages Impératifs d’Ordre Supérieur et Interprétation Abstraite*. PhD thesis, École Polytechnique, Paris, 1992. In French.
- [11] F. Bourdoncle. Efficient chaotic iteration strategies with widenings. In *Proceedings of the International Conference on Formal Methods in Programming and their Applications*, volume 735 of LNCS, pages 128–141. Springer, 1993.
- [12] S. P. Boyd and L. Vandenberghe. *Convex optimization*. Cambridge University Press, 2004. ISBN 0521833787.
- [13] M. A. Colón, S. Sankaranarayanan, and H. Sipma. Linear invariant generation using non-linear constraint solving. In W. A. Hunt, Jr. and F. Somenzi, editors, *Computer Aided Verification (CAV)*, number 2725 in LNCS, pages 420–433. Springer, 2003. ISBN 3-540-40524-0. doi: 10.1007/b11831.
- [14] A. Costan, S. Gaubert, É. Goubault, M. Martel, and S. Putot. A policy iteration algorithm for computing fixed points in static analysis of programs. In K. Etessami and S. K. Rajamani, editors, *Computer Aided Verification (CAV)*, number 4590 in LNCS, pages 462–475. Springer, 2005. ISBN 3-540-27231-3. doi: 10.1007/11513988_46.
- [15] P. Cousot. Proving program invariance and termination by parametric abstraction, Lagrangian relaxation and semidefinite programming. In R. Cousot, editor, *Verification, Model Checking and Abstract Interpretation (VMCAI)*, number 3385 in LNCS, pages 1–24. Springer, 2005. ISBN 3-540-24297-X. doi: 10.1007/b105073.
- [16] P. Cousot. *Méthodes itératives de construction et d’approximation de points fixes d’opérateurs monotones sur un treillis, analyse sémantique des programmes*. State doctorate thesis, Université scientifique et médicale de Grenoble & Institut national polytechnique de Grenoble, 1978. <http://tel.archives-ouvertes.fr/tel-00288657/en/>. In French.
- [17] P. Cousot and R. Cousot. Static determination of dynamic properties of programs. In *Proceedings of the Second International Symposium on Programming (1976)*, pages 106–130, Paris, 1977. Dunod. ISBN 2-04-005185-6. Also known as *Actes du deuxième colloque international sur la programmation*.
- [18] P. Cousot and R. Cousot. Abstract interpretation frameworks. *J. of Logic and Computation*, pages 511–547, Aug. 1992. ISSN 0955-792X. doi: 10.1093/logcom/2.4.511.

- [19] P. Cousot and N. Halbwachs. Automatic discovery of linear restraints among variables of a program. In *Principles of Programming Languages (POPL)*, pages 84–96. ACM, 1978. doi: 10.1145/512760.512770.
- [20] G. Dantzig. *Linear Programming and Extensions*. Princeton University Press, 1998. ISBN 0-691-05913-6.
- [21] B. Dutertre and L. de Moura. The Yices SMT solver. Tool paper, August 2006. <http://yices.cs1.sri.com/tool-paper.pdf>.
- [22] B. Dutertre and L. M. de Moura. A fast linear-arithmetic solver for DPLL(T). In T. Ball and R. B. Jones, editors, *CAV*, volume 4144 of *LNCS*, pages 81–94. Springer, 2006. ISBN 3-540-37406-X.
- [23] D. G. Espinoza. *On Linear Programming, Integer Programming and Cutting Planes*. PhD thesis, School of Industrial and Systems engineering, Georgia Institute of Technology, May 2006.
- [24] S. Gaubert, E. Goubault, A. Taly, and S. Zennou. Static analysis by policy iteration on relational domains. In R. D. Nicola, editor, *ESOP*, volume 4421 of *LNCS*, pages 237–252. Springer, 2007. ISBN 978-3-540-71314-2.
- [25] T. Gawlitza and D. Monniaux. Improving strategies via SMT solving. In G. Barthe, editor, *ESOP*, number 6602 in *LNCS*, pages 236–255. Springer, 2011. ISBN 978-3-642-19717-8. doi: 10.1007/978-3-642-19718-5_13.
- [26] T. Gawlitza and H. Seidl. Precise relational invariants through strategy iteration. In J. Duparc and T. A. Henzinger, editors, *CSL*, volume 4646 of *LNCS*, pages 23–40. Springer, 2007. ISBN 978-3-540-74914-1.
- [27] T. Gawlitza and H. Seidl. Precise interval analysis vs. parity games. In J. Cuéllar, T. S. E. Maibaum, and K. Sere, editors, *FM*, volume 5014 of *LNCS*, pages 342–357. Springer, 2008. ISBN 978-3-540-68235-6.
- [28] T. M. Gawlitza and H. Seidl. Solving systems of rational equations through strategy iteration. *ACM TOPLAS*, 33(3):11, 2011. doi: 10.1145/1961204.1961207.
- [29] T. M. Gawlitza and H. Seidl. Numerical invariants through convex relaxation and max-strategy iteration. <http://arxiv.org/abs/1204.1147>, 2012.
- [30] T. M. Gawlitza and H. Seidl. Computing relaxed abstract semantics w.r.t. quadratic zones precisely. In *Static analysis (SAS)*, *LNCS*. Springer, 2010.
- [31] T. M. Gawlitza and H. Seidl. Numerical Invariants through Convex Relaxation and Max-Strategy Iteration. Technical report, The University of Sydney, 2011.
- [32] L. Gonnord. *Accélération abstraite pour l’amélioration de la précision en analyse des relations linéaires*. PhD thesis, Université Joseph Fourier, Oct. 2007. <http://tel.archives-ouvertes.fr/tel-00196899/en/>.
- [33] L. Gonnord and N. Halbwachs. Combining widening and acceleration in linear relation analysis. In K. Yi, editor, *Static analysis (SAS)*, volume 4134 of *LNCS*, pages 144–160. Springer, 2006. ISBN 3-540-37756-5.
- [34] D. Gopan and T. W. Reps. Lookahead widening. In T. Ball and R. B. Jones, editors, *Computer Aided Verification (CAV)*, volume 4144 of *LNCS*, pages 452–466. Springer, 2006. ISBN 3-540-37406-X. doi: 10.1007/11817963.41.
- [35] D. Gopan and T. W. Reps. Guided static analysis. In H. R. Nielson and G. Filé, editors, *Static analysis (SAS)*, volume 4634 of *LNCS*, pages 349–365. Springer, 2007. ISBN 978-3-540-74060-5. doi: 10.1007/978-3-540-74061-2_22.
- [36] E. Goubault, S. L. Roux, J. Leconte, L. Liberti, and F. Marinelli. Static analysis by abstract interpretation: A mathematical programming approach. *Electronic notes in*

- theoretical computer science*, 267(1):73–87, 2010. doi: 10.1016/j.entcs.2010.09.007.
- [37] N. Halbwachs. *Détermination automatique de relations linéaires vérifiées par les variables d'un programme*. PhD thesis, Université scientifique et médicale de Grenoble and Institut National Polytechnique de Grenoble, 1979. <http://tel.archives-ouvertes.fr/tel-00288805/en/>.
- [38] N. Halbwachs and J. Henry. When the decreasing sequence fails. In *Static analysis (SAS)*, 2012. To appear.
- [39] N. Halbwachs, Y.-E. Proy, and P. Roumanoff. Verification of real-time systems using linear relation analysis. *Formal Methods in System Design*, 11(2):157–185, August 1997.
- [40] J. Henry, D. Monniaux, and M. Moy. Succinct representations for abstract interpretation. In *Static analysis (SAS)*, LNCS. Springer, 2012. To appear.
- [41] L. G. Khachiyan. A Polynomial Algorithm in Linear Programming. *Soviet Mathematics Doklady*, 20:191–194, 1979.
- [42] D. Kroening and O. Strichman. *Decision procedures*. Springer, 2008. ISBN 978-3-540-74104-6.
- [43] J. Leroux and G. Sutre. Accelerated data-flow analysis. In *Static Analysis (SAS)*, volume 4634 of *LNCS*, pages 184–199. Springer, 2007. doi: 10.1007/s10009-008-0064-3.
- [44] A. Miné. Symbolic methods to enhance the precision of numerical abstract domains. In E. A. Emerson and K. S. Namjoshi, editors, *VMCAI*, volume 3855 of *LNCS*, pages 348–363. Springer, 2006. ISBN 3-540-31139-4. doi: 10.1007/11609773_23.
- [45] A. Miné. The octagon abstract domain. *Higher-Order and Symbolic Computation*, 19(1):31–100, 2006. doi: 10.1007/s10990-006-8609-1.
- [46] A. Miné. *Domaines numériques abstraits faiblement relationnels*. PhD thesis, École polytechnique, 2004.
- [47] D. Monniaux. Quantifier elimination by lazy model enumeration. In *Computer-aided verification (CAV)*, number 6174 in *LNCS*, pages 585–599. Springer, 2010. ISBN 3642142958. doi: 10.1007/978-3-642-14295-6_51.
- [48] D. Monniaux. Automatic modular abstractions for template numerical constraints. *Logical Methods in Computer Science*, June 2010. ISSN 1860-5974. doi: 10.2168/LMCS-6(3:4)2010.
- [49] D. Monniaux and L. Gonnord. Using bounded model checking to focus fixpoint iterations. In E. Yahav, editor, *Static analysis (SAS)*, volume 6887 of *LNCS*, pages 369–385. Springer, 2011. ISBN 978-3-642-23701-0. doi: 10.1007/978-3-642-23702-7_27.
- [50] C. H. Papadimitriou. *Computational Complexity*. Addison-Wesley, 1994. ISBN 0-201-53082-1.
- [51] C. H. Papadimitriou and M. Yannakakis. A note on succinct representations of graphs. *Information and control*, 71:181–185, Dec. 1986. ISSN 0019-9958.
- [52] H. G. Rice. Classes of recursively enumerable sets and their decision problems. *Transactions of the American Mathematical Society*, 74(2):358–366, March 1953.
- [53] X. Rival and L. Mauborgne. The trace partitioning abstract domain. *Transactions on Programming Languages and Systems (TOPLAS)*, 29(5):26, 2007. ISSN 0164-0925. doi: 10.1145/1275497.1275501.
- [54] H. Rogers, Jr. *Theory of recursive functions and effective computability*. MIT Press, 1987. ISBN 0-262-68052-1.
- [55] S. Sankaranarayanan, H. Sipma, and Z. Manna. Constraint-based linear-relations analysis. In *Static Analysis (SAS)*, number 3148 in *LNCS*, pages 53–68. Springer, 2004.

- doi: 10.1007/b99688.
- [56] S. Sankaranarayanan, H. B. Sipma, and Z. Manna. Scalable analysis of linear systems using mathematical programming. In R. Cousot, editor, *VMCAI*, volume 3385 of *LNCS*, pages 25–41. Springer, 2005. ISBN 3-540-24297-X.
 - [57] A. Schrijver. *Theory of linear and integer programming*. John Wiley & Sons, Inc., New York, NY, USA, 1986. ISBN 0-471-90854-1.
 - [58] R. Sebastiani and S. Tomasi. Optimization in SMT with $\mathcal{LA}(\mathbb{Q})$ cost functions. In B. Gramlich, D. Miller, and U. Sattler, editors, *IJCAR*, volume 7364 of *LNCS*, pages 484–498. Springer, 2012. ISBN 978-3-642-31364-6. doi: 10.1007/978-3-642-31365-3_38.
 - [59] A. Shamir. A linear time algorithm for finding minimum cutsets in reducible graphs. *SIAM J. Comput.*, 8(4):645–655, 1979. ISSN 0097-5397. doi: 10.1137/0208051.
 - [60] J. Souyris and D. Delmas. Experimental assessment of Astrée on safety-critical avionics software. In F. Saglietti and N. Oster, editors, *SAFEComp*, volume 4680 of *LNCS*, pages 479–490. Springer, 2007. ISBN 978-3-540-75100-7. doi: 10.1007/978-3-540-75101-4_45.
 - [61] L. J. Stockmeyer. The polynomial-time hierarchy. *Theoretical Computer Science*, 3(1): 1–22, Oct. 1976. doi: 10.1016/0304-3975(76)90061-X.
 - [62] A. Tarski. A lattice-theoretical fixpoint theorem and its applications. *Pac. J. Math.*, 5:285–309, 1955. <http://projecteuclid.org/euclid.pjm/1103044538>.
 - [63] C. Wrathall. Complete sets and the polynomial-time hierarchy. *Theor. Comput. Sci.*, 3(1):23–33, 1976. ISSN 0304-3975. doi: 10.1016/0304-3975(76)90062-1.