

SubPolyhedra: A family of numerical abstract domains for the (more) scalable inference of linear inequalities

Vincent Laviron¹, Francesco Logozzo²

¹ École Normale Supérieure, 45, rue d’Ulm, Paris (France)

e-mail: Vincent.Laviron@ens.fr

² Microsoft Research, Redmond, WA (USA)

e-mail: logozzo@microsoft.com

The date of receipt and acceptance will be inserted by the editor

Abstract. We introduce Subpolyhedra (SubPoly) a new family of numerical abstract domains to infer and propagate linear inequalities. The key insight is that the reduced product of linear equalities and intervals produces powerful yet scalable analyses. Abstract domains in SubPoly are as expressive as Polyhedra, but they drop some of the deductive power to achieve scalability. The cost/precision ratio of abstract domains in the SubPoly family can be finely tuned according to: (i) the precision one wants to retain at join points!; and (ii) the algorithm used to infer tight bounds on intervals.

We implemented SubPoly on the top of Clousot, a generic abstract interpreter for .Net. Clousot with SubPoly analyzes very large and complex code bases in few minutes. SubPoly can efficiently capture linear inequalities among hundreds of variables, a result well-beyond state-of-the-art implementations of Polyhedra.

1 Introduction

The goal of an abstract interpretation-based static analyzer is to statically infer properties of the execution of a program that can be used to check its specification. The specification usually includes the absence of runtime exceptions (division by zero, integer overflow, array index out of bounds ...) and programmer annotations in the form of preconditions, postconditions, object invariants and assertions (“contracts” [30,2]). Proving that a piece of code satisfies its specification often requires discovering numerical invariants on program variables.

The concept of abstract domain is central in the design and the implementation of a static analyzer [10]. Abstract domains capture the properties of interest on programs. In particular *numerical* abstract domains are used to infer numerical relationships among program

```
class StringBuilder {
    int m_ChunkLength; char[] m_ChunkChars;
    // ...
    public void Append(int wb, int count) {
        Contract.Requires(wb >= 2 * count);
        if (count + m_ChunkLength > m_ChunkChars.Length)
    (*) CopyChars(wb, m_ChunkChars.Length - m_ChunkLength);
        // ... }

    private void CopyChars(int wb, int len) {
        Contract.Requires(wb >= 2 * len);
        // ...
    }
}
```

Fig. 1. An example extracted from mscorlib.dll. The function `Contract.Requires(...)` expresses method preconditions. Proving the precondition of `CopyChars` requires propagating an invariant involving three variables and non-unary coefficients.

variables. Cousot and Halbwachs introduced the Polyhedra numerical abstract domain (Poly) in [14]. Poly infers all the linear equalities on the program variables. The application and scalability of Poly has been severely limited by its performance which is worst-case exponential (easily attained in practice). To overcome this shortcoming and to achieve scalability, new numerical abstract domains have been designed moving in two orthogonal directions: either only considering inequalities of a particular shape (weakly relational domains) or fixing *ahead* of the analysis the maximum number of linear inequalities to be considered (bounded domains). The first class includes Octagons (which capture properties in the form $\pm x \pm y \leq c$) [31], TVPI ($a \cdot x + b \cdot y \leq c$) [38], Pentagons ($x \leq y \wedge a \leq x \leq b$) [29], Stripes ($x + a \cdot (y + z) > b$) [17] and Octahedra ($\pm x_0 \cdots \pm x_n \leq c$) [8]. The latter includes constraint template matrices (which capture at most m linear inequalities) [36,21] and methods to generate polynomial invariants *e.g.* [33,34,24].

Although impressive results have been achieved using weakly relational and bounded abstract domains, we experienced situations where the full *expressive* power of Poly is required. As an example, let us consider the code snippet of Fig. 1, extracted from `mscorlib.dll`, the main library of the .Net framework. Checking the precondition at the call site (*) involves:

- (i) *propagating* the given constraints:

$$wb \geq 2 \cdot \text{count}$$

$$\text{count} + \text{m_ChunkLength} > \text{m_ChunkChars.Length}$$

- (ii) *deducing* the precondition for `CopyChars`:

$$wb \geq 2 \cdot (\text{m_ChunkChars.Length} - \text{m_ChunkLength})$$

The aforementioned weakly relational domains cannot be used to check the precondition: Octahedra do not capture the first constraint (it involves a constraint with a non-unary coefficient); TVPI do not propagate the second constraint (it involves three variables); Pentagons and Octagons cannot represent any of the constraints; Stripes can propagate both constraints, but because of the incomplete closure it cannot deduce the precondition. Bounded domains does the job, provided we fix *before* the analysis the template for the constraints. This is inadequate for our purposes: The analysis of a *single* method in `mscorlib.dll` may involve hundreds of constraints, whose shape cannot be fixed ahead of the analysis, *e.g.* by a textual inspection. Poly easily propagates the constraints. However, in the general case the price to pay for using Poly is too elevated: the analysis will be limited to few tens of variables.

1.1 Subpolyhedra

We propose a new family of numerical abstract domains, Subpolyhedra (**SubPoly**), which has the same *expressive* power as Poly, but it drops some inference power to achieve scalability: **SubPoly** exactly represents and propagates linear inequalities containing hundreds of variables and constraints. **SubPoly** is based on the fundamental insight that the reduced product of linear equalities, `LinEq` [22], and intervals, `Intv` [10], can produce very powerful yet efficient program analyses. **SubPoly** can represent linear inequalities using slack variables, *e.g.* $wb \geq 2 \cdot \text{count}$ is represented in **SubPoly** by $wb - 2 \cdot \text{count} = \beta \wedge \beta \in [0, +\infty]$. As a consequence, **SubPoly** easily proves that the precondition for `CopyChars` is satisfied at the call site (*). In general the join of **SubPoly** is less precise than the one on Poly, so that it may not infer *all* the linear inequalities. The reason for that is that the pairwise join on `LinEq` and `Intv` is in general less precise than the join on Poly. To mitigate this loss of precision, we introduce a technique called hints [25] which enable recovering some of the precision. This technique is not

```
void Foo(int i, int j) {
    int x = i, y = j;
    if (x <= 0) return;
    while (x > 0) { x--; y--; }
    if (y == 0)
        Assert(i == j);
}
```

Fig. 2. An example from [35]. **SubPoly** infers the loop invariant $x - i = y - j \wedge x \geq 0$, propagates it and prove the assertion.

limited to **SubPoly**, and indeed we show that several existing refinement techniques can be seen as a particular case of hints.

The cardinal operation for **SubPoly** is the join, which computes a compact yet precise upper approximation of two incoming abstract states. The join of **SubPoly** is parameterized by the: (i) the reduction algorithm, which propagates the information between `LinEq` and `Intv`; and (ii) the hints, which recover information lost at join points. Every instantiation of the (reduction, hints) produces a new abstract domain in the **SubPoly** family, allowing the fine tuning of the cost/precision ratio. The most imprecise yet fast abstract domain in the **SubPoly** family is the one in which the reduction is the simple identity (no interval is refined) and the no hints are used. The most precise yet expensive abstract domain is one where the reduction is a complete linear programming algorithm and the hints are the usual Poly join.

1.2 Reduction

Let us consider the example in Fig. 2, taken from [35]. The program contains operations and predicates that can be exactly represented with Octagons. Proving that the assertion is not violated requires discovering the loop invariant $x - y = i - j \wedge x \geq 0$. The loop invariant cannot be fully represented in Octagons: it involves a relation on four variables. Bounded numerical domains are unlikely to help here as there is no way to syntactically figure out the required template. The `LinEq` component of **SubPoly** infers the relation $x - y = i - j$. The `Intv` component of **SubPoly** infers the loop invariant $x \in [0, +\infty]$, which in conjunction with the negation of the guard implies that $x \in [0, 0]$. The simplification of **SubPoly** propagates the interval, refining the linear constraint to $y = j - i$. This is enough to prove the assertion (in conjunction with the if-statement guard). It is worth noting that unlike [35] **SubPoly** does not require any hypothesis on the order of variables to prove the assertion.

1.3 Join and Hints

Let us consider the code in Fig. 3, taken from [19]. The loop invariant required to prove that the assertion is unreachable (and hence that the program is correct) is

```

int x = 0, y = 0, w = 0, z = 0;
while (...) {
  if (...) { x++; y += 100; }
  else if (...) { if (x >= 4) { x++; y++;; } }
  else if (y > 10 * w && z >= 100 * x) { y = -y; }

  w++; z += 10;
}
if (x >= 4 && y <= 2) Assert(false);

```

Fig. 3. An example from [19]. SubPoly infers the loop invariant $x \leq y \leq 100 \cdot x \wedge z = 10 \cdot w$, propagates it out of the loop, and proves that the assertion is unreachable.

$x \leq y \leq 100 \cdot x \wedge z = 10 \cdot w$. Without hints, SubPoly can only infer $z = 10 \cdot w$. *Template* hints, inspired by [36], are used to recover linear inequalities that are dropped by the imprecision of the join: In the example the template is $x - y \leq b$, and the analysis automatically figures out that $b = 0$. *Planar Convex hull* hints, inspired by [38], are used to introduce at join points linear inequalities derived by a planar convex hull: In the example it helps the analysis figure out that $y \leq 100 \cdot x$. It is worth noting that SubPoly does not need any of the techniques of [19] to infer the loop invariant.

2 Abstract Interpretation

2.1 Abstract domains

We assume the concrete domain to be the complete Boolean lattice of environments, $C = \langle \mathcal{P}(\Sigma), \subseteq, \emptyset, \Sigma, \cup, \cap \rangle$, where $\Sigma = [\text{Vars} \rightarrow \mathbb{Z}]$. An abstract domain A is a tuple $\langle \bar{D}, \gamma, \bar{\sqsubseteq}, \bar{\perp}, \bar{\top}, \bar{\sqcup}, \bar{\sqcap}, \bar{\nabla}, \rho, \sigma \rangle$. The set of abstract elements \bar{D} is related to the concrete domain by a *monotonic* concretization function $\gamma \in [\bar{D} \rightarrow C]$. With an abuse of notation, we will not distinguish between an abstract domain and the set of its elements. The approximation order is $\bar{\sqsubseteq}$ is a sound approximation of the concrete order:

$$\forall \bar{d}_0, \bar{d}_1 \in \bar{D}. \bar{d}_0 \bar{\sqsubseteq} \bar{d}_1 \implies \gamma(\bar{d}_0) \subseteq \gamma(\bar{d}_1)$$

The smallest element is $\bar{\perp}$, the largest element is $\bar{\top}$. The join operator $\bar{\sqcup}$ satisfies:

$$\forall \bar{d}_0, \bar{d}_1 \in \bar{D}. \bar{d}_0 \bar{\sqsubseteq} \bar{d}_0 \bar{\sqcup} \bar{d}_1 \wedge \bar{d}_1 \bar{\sqsubseteq} \bar{d}_0 \bar{\sqcup} \bar{d}_1$$

The meet operator $\bar{\sqcap}$ satisfies:

$$\forall \bar{d}_0, \bar{d}_1 \in \bar{D}. \bar{d}_0 \bar{\sqcap} \bar{d}_1 \bar{\sqsubseteq} \bar{d}_0 \wedge \bar{d}_0 \bar{\sqcap} \bar{d}_1 \bar{\sqsubseteq} \bar{d}_1$$

The widening $\bar{\nabla}$ ensures the convergence of the fixpoint iterations, *i.e.* it satisfies:

- (i) $\forall \bar{d}_0, \bar{d}_1 \in \bar{D}. \bar{d}_0 \bar{\sqsubseteq} \bar{d}_0 \bar{\nabla} \bar{d}_1 \wedge \bar{d}_1 \bar{\sqsubseteq} \bar{d}_0 \bar{\nabla} \bar{d}_1$
- (ii) for each sequence of abstract elements $\bar{d}_0, \bar{d}_1, \dots, \bar{d}_k$ the sequence defined by:
 $\bar{d}_0^\nabla = \bar{d}_0, \bar{d}_1^\nabla = \bar{d}_0 \bar{\nabla} \bar{d}_1 \dots \bar{d}_k^\nabla = \bar{d}_{k-1}^\nabla \bar{\nabla} \bar{d}_k$
 is ultimately stationary.

In general, we do not require abstract elements to be in some canonical or closed form, *i.e.* there may exist $\bar{d}_0, \bar{d}_1 \in \bar{D}$, such that $\bar{d}_0 \neq \bar{d}_1$, but $\gamma(\bar{d}_0) = \gamma(\bar{d}_1)$. The *reduction* operator $\rho \in [\bar{D} \rightarrow \bar{D}]$ puts an abstract element into a (pseudo-)canonical form without adding or losing any information: $\forall \bar{d}. \gamma(\rho(\bar{d})) = \gamma(\bar{d}) \wedge \rho(\bar{d}) \bar{\sqsubseteq} \bar{d}$. We do not require ρ to be idempotent. The *simplification* operator $\sigma \in [\bar{D} \rightarrow \bar{D}]$ removes redundancies in an abstract state. It may introduce some loss of precision: $\forall \bar{d}. \gamma(\bar{d}) \subseteq \gamma(\sigma(\bar{d}))$.

In most of the literature, reduction and simplification are not given the status of lattice operation. However, several domains use internally some specific operations that gives a more adapted representation of the abstract state, for a given operation. For instance, there is an operation on Octagons that is called *closure*, and which has the properties of a reduction operator. We believe that this is general enough to warrant adding two operators to the standard abstract domain definition. Of course, the identity is always a reduction operator and a simplification operator, so it can be defined even for domains which have no corresponding specific operation. Those operators are particularly important when the abstract elements considered are representations of mathematical objects, such that some objects have multiple equivalent representations.

New abstract domains can be systematically derived by cartesian composition or functional lifting [11]. Following [9], we use the dot-notation to denote point wise or functional extensions.

2.2 Transfer functions

It is common practice for the implementation of an abstract domain A to provide three abstract transfer functions: one for the assignment, one for the handling of tests, and one to perform abstract checking. The abstract transfer function for assignment, $A.\text{assign}$, is an over-approximation of the states reached after the concrete assignment ($\mathbb{E}[\mathbb{E}](\sigma)$ denotes the evaluation of the expression E in the state σ):

$$\forall \mathbf{x}, \mathbf{e}. \forall \bar{\mathbf{e}} \in A. \{ \sigma[\mathbf{x} \mapsto v] \mid \sigma \in \gamma(\bar{\mathbf{e}}), \mathbb{E}[\mathbf{E}](\sigma) = v \} \subseteq \gamma(A.\text{assign}(\bar{\mathbf{e}}, \mathbf{x}, \mathbf{E}))$$

The test abstract transfer function, $A.\text{test}$, filters the input states ($\mathbb{B}[\mathbb{B}](\sigma)$ denotes the evaluation of a Boolean expression B in the state σ):

$$\forall \bar{\mathbf{e}}. \forall \bar{\mathbf{e}} \in A. \{ \sigma \in \gamma(\bar{\mathbf{e}}) \mid \mathbb{B}[\mathbb{B}](\sigma) = \text{true} \} \subseteq \gamma(A.\text{test}(\bar{\mathbf{e}}, B))$$

The abstract checking $A.\text{check}$ verifies if an assertion A holds in an abstract state $\bar{\mathbf{e}}$. It has four possible outcomes: *true* meaning that A holds in all the concrete states $\gamma(\bar{\mathbf{e}})$; *false*, meaning that $\neg A$ holds in all the concrete states $\gamma(\bar{\mathbf{e}})$; *bottom*, meaning that the assertion is unreachable; *top*, meaning that the validity of A cannot be decided in $\gamma(\bar{\mathbf{e}})$. Formally, $A.\text{check}$ satisfies $\forall A. \forall \bar{\mathbf{e}} \in A$:

Order:	$[a_1, b_1] \sqsubseteq_{\text{Intv}} [a_2, b_2] \iff a_1 \geq a_2 \wedge b_1 \leq b_2$
Bottom:	$[a, b] = \perp_{\text{Intv}} \iff a > b$
Top:	$[a, b] = \top_{\text{Intv}} \iff a = -\infty \wedge b = +\infty$
Join:	$[a_1, b_1] \sqcup_{\text{Intv}} [a_2, b_2] = [\min(a_1, a_2), \max(b_1, b_2)]$
Meet:	$[a_1, b_1] \sqcap_{\text{Intv}} [a_2, b_2] = [\max(a_1, a_2), \min(b_1, b_2)]$
Widening:	$[a_1, b_1] \nabla_{\text{Intv}} [a_2, b_2] =$ $\quad \text{if } a_1 > a_2 \text{ then } a_2 \text{ else } -\infty,$ $\quad \text{if } b_1 < b_2 \text{ then } b_2 \text{ else } +\infty$

Fig. 4. Lattice operations over single intervals

$\text{A.check}(\mathbf{A}, \bar{e}) = v \Rightarrow \forall \sigma \in \gamma(\bar{e}). \mathbb{B}[\mathbf{A}](\sigma) = v, v \in \{\text{true}, \text{false}\}$
 $\text{A.check}(\mathbf{A}, \bar{e}) = \text{bot} \Rightarrow \gamma(\bar{e}) = \emptyset$
 $\text{A.check}(\mathbf{A}, \bar{e}) = \text{top} \Rightarrow \exists \sigma_0, \sigma_1 \in \gamma(\bar{e}). \mathbb{B}[\mathbf{A}](\sigma_0) \neq \mathbb{B}[\mathbf{A}](\sigma_1)$

2.3 Intervals

The abstract domain of interval environments is $\langle \text{Intv}, \gamma_{\text{Intv}}, \perp_{\text{Intv}}, \top_{\text{Intv}}, \sqcup_{\text{Intv}}, \sqcap_{\text{Intv}}, \nabla_{\text{Intv}} \rangle$. The abstract elements are maps from program variables to open intervals. The concretization of an interval environment \bar{i} is

$$\gamma_{\text{Intv}}(\bar{i}) = \{s \in \Sigma \mid \forall \mathbf{x} \in \text{dom}(\bar{i}). \bar{i}(\mathbf{x}) = [a, b] \wedge a \leq s(\mathbf{x}) \leq b\}.$$

The lattice operations are the functional extension of those in Fig. 4. The reduction and the simplification for intervals are the identity function. All the domain operations can be implemented in linear time.

2.4 Linear Equalities

The abstract domain of linear *equalities* is $\langle \text{LinEq}, \gamma_{\text{LinEq}}, \perp_{\text{LinEq}}, \top_{\text{LinEq}}, \sqcup_{\text{LinEq}}, \sqcap_{\text{LinEq}} \rangle$. The elements are sets of linear equalities, their meaning is given by the set of concrete states which satisfy the constraints, *i.e.*

$$\gamma_{\text{LinEq}} = \lambda \bar{l}. \{s \in \Sigma \mid \forall (\sum a_i \cdot \mathbf{x}_i = b) \in \bar{l}. \sum a_i \cdot s(\mathbf{x}_i) = b\}.$$

The order is sub-space inclusion, the bottom is the empty space, the top is the whole space, the join is the smallest space which contains the two arguments, the meet is space intersection. LinEq satisfies the ascending chain condition, so that the join suffices to ensure analysis termination. The reduction and the simplification are just Gaussian elimination. The complexity of the domain operations is subsumed by the complexity of Gaussian elimination, which is cubic.

2.5 Polyhedra

The abstract domain of linear *inequalities* is $\langle \text{Poly}, \gamma_{\text{Poly}}, \perp_{\text{Poly}}, \top_{\text{Poly}}, \sqcup_{\text{Poly}}, \sqcap_{\text{Poly}}, \nabla_{\text{Poly}} \rangle$. The elements are sets of linear inequalities, the concretization is the set of concrete states which satisfy the constraints *i.e.*

$$\gamma_{\text{Poly}} = \lambda \bar{p}. \{s \in \Sigma \mid \forall (\sum a_i \cdot \mathbf{x}_i \leq b) \in \bar{p}. \sum a_i \cdot s(\mathbf{x}_i) \leq b\},$$

the order is the polyhedron inclusion, the bottom is the empty polyhedron, the top is the whole space, the join is the convex hull, the meet is just the union of the set of constraints, and the widening preserves the inequalities stable among two successive iterations. The reduction and the simplification respectively infers the set of generators and removes the redundant inequalities. The cost of the Poly operations is subsumed by the cost of the conversion between the algebraic representation (set of inequalities) and the geometric representation (set of generators) used in the implementation [1]. In fact, some operations require the algebraic representation (*e.g.* $\bar{\Pi}_{\text{Poly}}$), some require the geometrical representation (*e.g.* $\bar{\Pi}_{\text{Poly}}$), and some others require both (*e.g.* $\bar{\square}_{\text{Poly}}$). The conversion between the two representations is exponential in the number of variables, and it cannot be done better [23].

3 Subpolyhedra

We introduce the numerical abstract domain of Subpolyhedra, SubPoly . The main idea of SubPoly is to combine Intv and LinEq to capture complex linear *inequalities*. Slack variables are introduced to replace inequality constraints with equalities.

3.1 Variables

A variable $\mathbf{v} \in \mathbf{Vars}$ can either be a *program* variable ($\mathbf{x} \in \mathbf{Var}_P$) or a *slack* variable ($\beta \in \mathbf{Var}_S$). A slack variable β has an associated information, denoted by $\text{info}(\beta)$, which is a linear form $a_1 \cdot \mathbf{v}_1 + \dots + a_k \cdot \mathbf{v}_k$.

Let $\kappa \equiv \sum a_i \cdot \mathbf{x}_i + \sum b_j \cdot \beta_j = c$ be a linear equality: $\mathbf{s}_\kappa = \sum_{\mathbf{x}_i \in \mathbf{Var}_P} a_i \cdot \mathbf{x}_i$ denotes the partial sum of the monomials involving just program variables; $\mathbf{Var}_P(\kappa) = \{\mathbf{x}_i \mid a_i \cdot \mathbf{x}_i \in \kappa, a_i \neq 0\}$ and $\mathbf{Var}_S(\kappa) = \{\beta_j \mid b_j \cdot \beta_j \in \kappa, b_j \neq 0\}$ denote respectively the program variables and the slack variables in κ . The generalization to inequalities and sets of equalities and inequalities is straightforward.

3.2 Domain structure

The elements of SubPoly belong to the reduced product $\text{LinEq} \times \text{Intv}$ [11]. Inequalities are represented in SubPoly with slack variables:

$$\sum a_i \cdot \mathbf{x}_i \leq c \iff \sum a_i \cdot \mathbf{x}_i - c = \beta \wedge \beta \in [-\infty, 0]$$

(β is a fresh slack variable with the associated information $\text{info}(\beta) = \sum a_i \cdot \mathbf{x}_i$).

3.3 Concretization

An element of SubPoly can be interpreted as a polyhedron by projecting out the slack variables: $\gamma_S^{\text{Poly}} \in$

$[\text{SubPoly} \rightarrow \text{Poly}]$ is

$$\gamma_S^{\text{Poly}} = \lambda(\bar{l}; \bar{i}). \pi_{\text{vars}}(\bar{l} \cup \{a \leq v \leq b \mid \bar{i}(v) = [a, b]\}),$$

where π denotes the projection of variables in Poly . The concretization $\gamma_S \in [\text{SubPoly} \rightarrow \mathcal{P}(\Sigma)]$ is then $\gamma_S = \gamma_{\text{Poly}} \circ \gamma_S^{\text{Poly}}$.

3.4 Approximation Order

The order on SubPoly may be defined in terms of order over Poly . Given two subpolyhedra \bar{s}_0, \bar{s}_1 , the most precise order relation $\bar{\sqsubseteq}_S^*$ is

$$\bar{s}_0 \bar{\sqsubseteq}_S^* \bar{s}_1 \iff \gamma_S^{\text{Poly}}(\bar{s}_0) \bar{\sqsubseteq}_{\text{Poly}} \gamma_S^{\text{Poly}}(\bar{s}_1).$$

However, $\bar{\sqsubseteq}_S^*$ may be too expensive to compute: it involves mapping subpolyhedra in the dual representation of Poly . This can easily cause an exponential blow up. We define a weaker approximation order relation which first tries to find a renaming θ for the slack variables, and then checks the pairwise order. Formally:

$$\begin{aligned} \langle \bar{l}_0; \bar{i}_0 \rangle \bar{\sqsubseteq}_S \langle \bar{l}_1; \bar{i}_1 \rangle &\iff \\ \exists \theta. \text{Var}_S(\langle \bar{l}_0; \bar{i}_0 \rangle) &\xrightarrow{\text{inj}} \text{Var}_S(\langle \bar{l}_1; \bar{i}_1 \rangle). \\ \forall \beta \in \text{Var}_S(\langle \bar{l}_0; \bar{i}_0 \rangle). \text{info}(\beta) &= \text{info}(\theta(\beta)) \\ &\wedge \theta(\langle \bar{l}_0; \bar{i}_0 \rangle) \dot{\sqsubseteq} \langle \bar{l}_1; \bar{i}_1 \rangle. \end{aligned}$$

In general $\bar{\sqsubseteq}_S \subsetneq \bar{\sqsubseteq}_S^*$. In practice, $\bar{\sqsubseteq}_S$ is used to check if a fixpoint has been reached. A weaker order relation means that the analysis may perform some extra widening steps, which may introduce precision loss. However, we found the definition of $\bar{\sqsubseteq}_S$ satisfactory in our experience.

One other important consequence of using a weak approximation order is that we are not always able to tell whether two abstract elements are actually equivalent representations of the same geometric shape. This is why, unlike some other domains like Poly , in SubPoly the elements do not correspond to a geometrical shape, even up to equivalence; there are some elements that correspond to the same polyhedron, but are not comparable with our weak ordering.

3.5 Bottom

An element of SubPoly is equivalent to bottom if after a reduction one of the two components is bottom:

$$\bar{s} = \perp_S \iff \rho(\bar{s}) = (\bar{l}, \bar{i}) \wedge (\bar{i} = \dot{\perp}_{\text{Intv}} \vee \bar{l} = \perp_{\text{LinEq}}).$$

3.6 Top

An element of SubPoly is top if after the simplification both components are top:

$$\bar{s} = \top_S \iff \sigma(\bar{s}) = (\bar{l}, \bar{i}) \wedge \bar{i} = \dot{\top}_{\text{Intv}} \wedge \bar{l} = \top_{\text{LinEq}}.$$

Algorithm 1 The join \sqcup_S on Subpolyhedra

input $\langle \bar{l}_i; \bar{i}_i \rangle \in \text{SubPoly}$, $i \in \{0, 1\}$

let $\langle \bar{l}'_i; \bar{i}'_i \rangle = \langle \bar{l}_i; \bar{i}_i \rangle$

{Step 1. Propagate the information of the slack variables}

for all $\beta \in \text{Var}_S(\bar{l}'_i) \setminus \text{Var}_S(\bar{i}'_i)$ **do**

$$\langle \bar{l}'_i; \bar{i}'_i \rangle := \langle \bar{l}'_i \sqcap_{\text{LinEq}} \{\beta = \text{info}(\beta)\}; \bar{i}'_i \rangle$$

{Step 2. Perform the point-wise join on the saturated operands}

let $\langle \bar{l}_\sqcup; \bar{i}_\sqcup \rangle = \rho(\langle \bar{l}'_0; \bar{i}'_0 \rangle) \dot{\sqcup} \rho(\langle \bar{l}'_1; \bar{i}'_1 \rangle)$

{Step 3. Hints: Recover the lost information }

let D_i be the linear equalities dropped from \bar{l}'_i at the previous step

for all $\kappa \in D_i$ **do**

if κ contains no slack variable **then**

let $\bar{i}_{s_\kappa} = \llbracket \mathbf{s}_\kappa \rrbracket \langle \bar{l}'_i; \bar{i}'_i \rangle$

if $\bar{i}_{s_\kappa} \neq \top_{\text{Intv}}$ **then**

let β be a fresh slack variable

$$\langle \bar{l}_\sqcup; \bar{i}_\sqcup \rangle := \langle \bar{l}_\sqcup \sqcap_{\text{LinEq}} \{\beta = \kappa\}; \bar{i}_\sqcup \dot{\sqcap}_{\text{Intv}} \{\beta = \bar{i}_{s_\kappa} \sqcup_{\text{Intv}} [0, 0]\} \rangle$$

else if κ contains exactly one slack variable β **then**

let $\bar{i}_{s_\kappa} = \llbracket \mathbf{s}_\kappa \rrbracket \langle \bar{l}'_i; \bar{i}'_i \rangle$

if $\bar{i}_{s_\kappa} \neq \top_{\text{Intv}}$ **then**

$$\langle \bar{l}_\sqcup; \bar{i}_\sqcup \rangle := \langle \bar{l}_\sqcup \sqcap_{\text{LinEq}} \{\kappa\}; \bar{i}_\sqcup \dot{\sqcap}_{\text{Intv}} \{\beta = \bar{i}_{s_\kappa} \sqcup_{\text{Intv}} \bar{i}_i(\beta)\} \rangle$$

else $\{\kappa$ contains strictly more than one slack variable}

continue

return $\langle \bar{l}_\sqcup; \bar{i}_\sqcup \rangle$

3.7 Linear form evaluation

Let \mathbf{s} be a linear form: $\llbracket \mathbf{s} \rrbracket \in [\text{SubPoly} \rightarrow \text{Intv}]$ denotes the evaluation of \mathbf{s} in an element of SubPoly after that the reduction has inferred the tightest bounds: $\llbracket \sum a_i \cdot v_i \rrbracket \langle \bar{l}; \bar{i} \rangle = \text{let } \langle \bar{l}^*; \bar{i}^* \rangle = \rho(\langle \bar{l}; \bar{i} \rangle) \text{ in } \sum a_i \cdot \bar{i}^*(v_i)$.

3.8 Join

As with the order, one could define a most precise join operation by concretizing on Poly , doing the convex hull, then abstracting again to SubPoly . However, this is a very expensive operation, and the aim of SubPoly is to give faster, but potentially less precise, operations. So we define instead a specific join algorithm \sqcup_S in three steps. First, inject the information of the slack variables into the abstract elements. Second, perform the pairwise join on the saturated arguments. Third, add the constraints that are implied by the two operands of the join, but that were not preserved by the previous step. The join is defined by the Algorithm 1 (We let $\underline{0} = 1$, $\underline{1} = 0$). We illustrate it with examples.

Example 1 (Steps 1 & 2). Let us consider the code in Fig. 5(a). After the assumption, the abstract states on the left branch and the right branch are respectively:

```

if(...)
  { assume x - y <= 0; }
else
  { assume x - y <= 5; }
      (a)

if(...)
  { assume x == y; assume y <= z; }
else
  { assume x <= y; assume y == z; }
      (b)

```

Fig. 5. Examples illustrating the need for Step 1 in the join algorithm

$\bar{s}_0 = \langle x - y = \beta_0; \beta_0 \in [-\infty, 0] \rangle$ and $\bar{s}_1 = \langle x - y = \beta_1; \beta_1 \in [-\infty, 5] \rangle$. The information associated with the slack variables is $\text{info}(\beta_0) = \text{info}(\beta_1) = x - y$. At the join point we apply Algorithm 1. Step 1 refines the abstract states by introducing the information associated with the slack variables: $\bar{s}'_0 = \langle x - y = \beta_0 = \beta_1; \beta_0 \in [-\infty, 0] \rangle$ and $\bar{s}'_1 = \langle x - y = \beta_1 = \beta_0; \beta_1 \in [-\infty, 5] \rangle$. Step 2 requires the reduction of the operands. The interval for β_1 (resp. β_0) in \bar{s}'_0 (resp. \bar{s}'_1) is refined: $\rho(\bar{s}'_0) = \langle x - y = \beta_0 = \beta_1; \beta_0 \in [-\infty, 0], \beta_1 \in [-\infty, 0] \rangle$ and $\rho(\bar{s}'_1) = \langle x - y = \beta_1 = \beta_0; \beta_0 \in [-\infty, 5], \beta_1 \in [-\infty, 5] \rangle$. The pairwise join gets the expected invariant: $\bar{s}_\sqcup = \rho(\bar{s}'_0) \sqcup \rho(\bar{s}'_1) = \langle x - y = \beta_0 = \beta_1; \beta_0 \in [-\infty, 5], \beta_1 \in [-\infty, 5] \rangle$. \square

Example 2 (Non-trivial information for slack variables). Let us consider the code snippet in Fig. 5(b). The abstract states to be joined are $\langle x - y = 0, y - z = \beta_0; \beta_0 \in [-\infty, 0] \rangle$ and $\langle y - z = 0, x - y = \beta_1; \beta_1 \in [-\infty, 0] \rangle$. The associated information are $\text{info}(\beta_0) = y - z$ and $\text{info}(\beta_1) = x - y$. Step 1 allows to refine the abstract states with the slack variable information, and hence to infer that after the join $x \leq y$ and $y \leq z$. \square

The two examples above show the importance of introducing the information associated with slack variables in Step 1 and the reduction in Step 2. Without those, the relation between the slack variables and the program point where they were introduced would have been lost.

The join of LinEq is *precise* in that if a linear equality is implied by both operands, then it is implied by the result too. The same for the join of Intv . The pairwise join in $\text{LinEq} \times \text{Intv}$ may drop some inequalities. Some of those can be recovered by the refinement step. The next example illustrates it.

Example 3 (Step 3). Let us consider the code in Fig. 6(a). The analysis of the two branches of the conditional produces the abstract states: $\bar{s}_0 = \langle x - 3 \cdot y = 0; \top_{\text{Intv}} \rangle$ and $\bar{s}_1 = \langle x = 0, y = 1; x \in [0, 0], y \in [1, 1] \rangle$. The reduction ρ does not refine the states (we already have the tightest bounds). The point-wise join produces the abstract state \top_S . Step 3 identifies the dropped constraints: $D_0 = \{x - 3 \cdot y = 0\}$ and $D_1 = \{x = 0, y = 1\}$. The algorithm inspects them to check if the corresponding linear form

```

if(...) { assume x == 3 * y; }
else { x = 0; y = 1; }
      (a)

i := k;
while(...) i++;
assert i >= k;
      (b)

```

Fig. 6. Examples illustrating the need for the Step 3 in the join and the widening.

can be bounded by the “other” branch. The linear form in D_0 is also bounded in the right branch: $\llbracket x - 3 \cdot y \rrbracket(\bar{s}_1) = [-3, -3] (\neq \top_{\text{Intv}})$. Therefore it is meaningful to add a slack variable β corresponding to this linear form to the result. The linear formss of D_1 cannot be bounded on the left branch, so they are discarded. The abstract state after the join is then $\bar{s}_\sqcup = \langle x - y = \beta; \beta \in [-3, 0] \rangle$. \square

3.9 Meet

The meet \sqcap_S is simply the pairwise meet on $\text{LinEq} \times \text{Intv}$.

3.10 Widening

The algorithm for the widening is similar to the join, with the main differences that: (i) the information associated to slack variables is propagated only in one direction; (ii) only the right argument is saturated; and (iii) the recovery step is applied only to one of the operands. Those hypotheses avoid the well-known problems of interaction between reduction, refinement and convergence of the iterations [32, 37].

Example 4 (Refinement step for the widening). Let us consider the code snippet in Fig. 6(b). The entry state to the loop is $\bar{s}_0 = \langle i - k = 0; \top_{\text{Intv}} \rangle$. The state after one iteration is $\bar{s}_1 = \langle i - k = 1; \top_{\text{Intv}} \rangle$. We apply the widening operator. Step 1 does not refine the states as there are no slack variables. The pairwise widening of Step 2 lose all the information. Step 3 recovers the constraint $k \leq i$: $D_0 = \{i - k = 0\}$ contains no slack variables and $\llbracket i - k \rrbracket(\bar{s}_1) = [1, 1]$ so that $\bar{s}_\nabla = \langle i - k = \beta; \beta \in [0, +\infty] \rangle$. \square

Theorem 1 (Fixpoint convergence). *The operator defined in Algorithm 2 is a widening. Moreover, $\bar{\sqsubseteq}_S$ can be used to check that the fixpoint iterations eventually stabilize.*

Proof sketch. Algorithm 2 ensures that the number of linear equalities at any step is at most the number of equalities in the first step. So there exists a point from which no more slack variables will be added. Existing slack variables may be renamed to fresh ones to avoid conflicts. In the definition of $\bar{\sqsubseteq}_S$ the renaming θ takes

Algorithm 2 The widening ∇_S on Subpolyhedra

```
input  $\langle \bar{l}_i; \bar{i}_i \rangle \in \text{SubPoly}$ ,  $i \in \{0, 1\}$ 
let  $\langle \bar{l}'_i; \bar{i}'_i \rangle = \langle \bar{l}_i; \bar{i}_i \rangle$ 
{Step 1. Propagate the information of the slack variables}
for all  $\beta \in \text{Vars}(\bar{l}_0) \setminus \text{Vars}(\bar{l}_1)$  do
   $\langle \bar{l}'_0; \bar{i}'_0 \rangle := \langle \bar{l}'_0 \sqcap_{\text{LinEq}} \{\beta = \text{info}(\beta)\}; \bar{i}'_0 \rangle$ 
{Step 2. Perform the point-wise widening}
let  $\langle \bar{l}'_{\nabla}; \bar{i}'_{\nabla} \rangle = \langle \bar{l}'_0; \bar{i}'_0 \rangle \dot{\nabla} \rho(\langle \bar{l}'_1; \bar{i}'_1 \rangle)$ 
{Step 3. Recover the lost information}
let  $D_0$  be the linear equalities dropped from  $\bar{l}'_0$  at the previous step
for all  $\kappa \in D_0$  do
  if  $\kappa$  contains no slack variables then
    let  $\bar{i}_{s\kappa} = \llbracket \mathbf{s}_{\kappa} \rrbracket(\bar{l}'_1; \bar{i}'_1)$ 
    if  $\bar{i}_{s\kappa} \neq \top_{\text{Intv}}$  then
      let  $\beta$  be a fresh slack variable
       $\langle \bar{l}'_{\nabla}; \bar{i}'_{\nabla} \rangle := \langle \bar{l}'_{\nabla} \sqcap_{\text{LinEq}} \{\beta = \kappa\}; \bar{i}'_{\nabla} \sqcap_{\text{Intv}} \{\beta = [0, 0] \bar{\nabla} \bar{i}_{s\kappa} \} \rangle$ 
    else if  $\kappa$  contains exactly one slack variable  $\beta$  then
      let  $\bar{i}_{s\kappa} = \llbracket \mathbf{s}_{\kappa} \rrbracket(\bar{l}'_1; \bar{i}'_1)$ 
      if  $\bar{i}_{s\kappa} \neq \top_{\text{Intv}}$  then
         $\langle \bar{l}'_{\nabla}; \bar{i}'_{\nabla} \rangle := \langle \bar{l}'_{\nabla} \sqcap_{\text{LinEq}} \{\kappa\}; \bar{i}'_{\nabla} \sqcap_{\text{Intv}} \{\beta = \bar{i}_{s\kappa}(\mathbf{v}) \bar{\nabla} \bar{i}_{s\kappa} \} \rangle$ 
      else  $\{\kappa$  contains strictly more than one slack variable}
    continue
  return  $\langle \bar{l}'_{\nabla}; \bar{i}'_{\nabla} \rangle$ 
```

care of those. Up to the renaming, the widening is the pairwise widening, which is convergent and whose stability can be checked by the pairwise partial order. \square

4 Reduction for Subpolyhedra

The reduction in `SubPoly` infers tighter bounds on linear forms and hence on program variables. Reduction is cardinal to fine tuning the precision/cost ratio. We propose two reduction algorithms, one based on linear programming, ρ_{LP} , and the other on basis exploration, ρ_{BE} . Both of them have been implemented in `Clousot`, our abstract interpretation-based static analyzer for `.Net` [16].

4.1 Linear programming-based reduction

A linear programming problem is the problem of maximizing (or minimizing) a linear function subject to a finite number of linear constraints. We consider *upper bounding* linear problems (UBLP) [7], *i.e.* problems in the form (n is the number of variables, m is the number

of equations):

$$\begin{aligned} & \text{maximize} && c \cdot \mathbf{v}_k && k \in 1 \dots n, c \in \{-1, +1\} \\ & \text{subject to} && \sum_{j=1}^n a_{ij} \cdot \mathbf{v}_j = b_j && (i = 1, \dots, m) \\ & && \text{and } l_j \leq \mathbf{v}_j \leq u_j && (j = 1, \dots, n). \end{aligned}$$

The Linear programming-based reduction ρ_{LP} is trivially an instance of UBLP: To infer the tightest upper bound (resp. lower bound) on a variable \mathbf{v}_k in an element of `SubPoly` $\langle \bar{l}; \bar{i} \rangle$ instantiate UBLP with $c = 1$ (resp. $c = -1$) subject to the linear equalities \bar{l} and the numerical bounds \bar{i} .

UBLP can be solved in polynomial time [7]. However, polynomial time algorithms for UBLP do not perform well in practice. The Simplex method [15], exponential in the worst-case, in practice performs a lot better than other known linear programming algorithms [39]. The Simplex algorithm works by visiting the *feasible bases* (informally, the vertices) of the polyhedron associated with the constraints. At each step, the algorithm visits the adjacent basis (vertex) that maximizes the current value of the objective by the largest amount. The iteration strategy of the Simplex guarantees the convergence to a basis which exhibits the optimal value for the objective.

The advantages of using Simplex for ρ_{LP} are that: (i) it is well-studied and optimized; (ii) it is complete in \mathbb{R} , *i.e.* it finds the best solution over real numbers; and (iii) it guarantees that all the information is propagated at once: $\rho_{LP} \circ \rho_{LP} = \rho_{LP}$.

The drawbacks of using Simplex are that (i) the computation over machine floating point may introduce imprecision or unsoundness in the result; and (ii) the reduction ρ_{LP} requires to solve $2 \cdot n$ UBLP problems to find the lower bound and the upper bound for each of the n variables in an abstract state. We have observed (i) in our experiences (cf. Sect. 7). There exists methods to circumvent the problem at the price of extra computational cost, *e.g.* using arbitrary precision rationals, or a combination of machine floating arithmetic and precise arithmetic. Even if (i) is solved, we observed that (ii) dominates the cost of the reduction, in particular in presence of abstract states with a large number of variables: the $2 \cdot n$ UBLP problems are *disjoints* and there is no easy way to share the sequence of bases visited by the Simplex algorithm over the different runs of the algorithm for the same abstract state.

4.2 Basis exploration-based reduction

We have developed a new reduction ρ_{BE} , less subject to the drawbacks from floating point computation than ρ_{LP} , which enables a better tuning of the precision/cost ratio than the Simplex. The basic ideas are: (i) to fix *ahead* of time the bases we want to explore; and (ii)

Algorithm 3 The reduction algorithm ρ_{BE} , parametrized by the oracle δ

```

input  $\langle \bar{l}; \bar{i} \rangle \in \text{SubPoly}$ ,  $\delta \in \mathcal{P}(\{\zeta \mid \zeta \text{ is a basis change}\})$ 

Put  $\bar{l}$  into row echelon form. Call the result  $\bar{l}'$ 
let  $\langle \bar{l}^*, \bar{i}^* \rangle = \langle \bar{l}', \bar{i} \rangle$ 
for all  $\zeta \in \delta$  do
     $\bar{l}^* := \zeta(\bar{l}^*)$ 
    for all  $\mathbf{v}_k + a_{k+1} \cdot \mathbf{v}_{k+1} + \dots + a_n \cdot \mathbf{v}_n = b \in \bar{l}^*$  do
         $\bar{i}^* := \bar{i}^*[\mathbf{v}_k \mapsto \bar{i}^*(\mathbf{v}_k) \sqcap_{\text{Intv}} \llbracket b - a_{k+1} \cdot \mathbf{v}_{k+1} + \dots + a_n \cdot \mathbf{v}_n \rrbracket(\bar{i}^*)]$ 
    return  $\langle \bar{l}^*, \bar{i}^* \rangle$ 

```

to refine at each step the variable bounds. The reduction ρ_{BE} , parametrized by a set of changes of basis δ , is formalized by Algorithm 3. First, we put the initial set of linear constraints into triangular form (row echelon form). Then, we apply the basis changes in δ and we refine all the variables *in the basis*. With respect to ρ_{LP} , ρ_{BE} is faster: (i) the number of bases to explore is statically bounded; (ii) at each step, k variables may be refined at once.

In theory, ρ_{BE} is an abstraction of ρ_{LP} , in that it may not infer the *optimal* bounds on variables (it depends on the choice of δ). In practice, we found that ρ_{BE} is much more numerically stable and it can infer better bounds than ρ_{LP} . The reason is in the handling of numerical errors in the computation. Suppose we are seeking a (lower or upper) bound for a variable using the Simplex. If we detect a numerical error (*i.e.* a loss of precision in floating point computations or a huge coefficient in the exact arithmetic computation), the only sound solution is to stop the iterations, and return the current value of the objective function as the result. On the other hand, when we detect a numerical error in ρ_{BE} , we can just skip the current basis (abstraction), and move to the next one in δ .

4.2.1 Linear Explorer (δ_L)

The linear bases explorer is based on the empirical observation that in most cases, to infer the tightest bounds for some variable \mathbf{v}_0 , you need to have it in the basis while some other variable \mathbf{v}_1 is out of the basis. Following this, the linear explorer generates a sequence of bases δ_L with the property that for each unordered pair of distinct variables $\langle \mathbf{v}_0, \mathbf{v}_1 \rangle$, there exists $\zeta \in \delta_L$ such that \mathbf{v}_0 is in the basis and \mathbf{v}_1 is not. The sequence δ_L is defined as $\delta_L = \{\zeta_i \mid i \in [0, n], \mathbf{v}_i \dots \mathbf{v}_{(i+m-1) \bmod n} \text{ are in basis for } \zeta_i\}$.

Example 5. (Reduction with the linear explorer) Let the initial state be $\bar{s} = \langle \mathbf{v}_0 + \mathbf{v}_2 + \mathbf{v}_3 = 1, \mathbf{v}_1 + \mathbf{v}_2 - \mathbf{v}_3 = 0; \mathbf{v}_0 \in [0, 2], \mathbf{v}_1 \in [0, 3] \rangle$, so that $\delta_L = \{\{\mathbf{v}_0, \mathbf{v}_1\}, \{\mathbf{v}_1, \mathbf{v}_2\}, \{\mathbf{v}_2, \mathbf{v}_3\}, \{\mathbf{v}_3, \mathbf{v}_0\}\}$. The reduction $\rho_{BE}(\bar{s})$ contains the tightest bounds for $\mathbf{v}_2, \mathbf{v}_3$: $\langle \mathbf{v}_2 + \frac{1}{2} \cdot \mathbf{v}_0 + \frac{1}{2} \cdot \mathbf{v}_1 = 0, \mathbf{v}_3 + \frac{1}{2} \cdot \mathbf{v}_0 - \frac{1}{2} \cdot \mathbf{v}_1 = 0; \mathbf{v}_0 \in [0, 2], \mathbf{v}_1 \in [0, 3], \mathbf{v}_2 \in [-\frac{5}{2}, 0], \mathbf{v}_3 \in [-1, \frac{3}{2}] \rangle$. \square

```

void AbsEl(int x)
{ if(...) x == -1;
  else    x = 1;

  assert  x != 0;
}
(a)

void Transfer(int x, y)
{ assume 2 <= x <= 3;
  assume -1 <= y <= 1;

  int z = (x + y) * y;

  assert -2 <= z;
}
(b)

```

Fig. 7. Examples of orthogonal losses of precision in abstract interpretations: (a) a convex domain cannot represent $x \neq 0$; and (b) a compositional transfer function does not infer the tightest lower bound for z

Properties of δ_L are that: (i) each variable appears exactly m times in the basis; (ii) it can be implemented efficiently as the basis change from ζ_i to ζ_{i+1} , $i \in [0, n-1]$ requires just one variable swap; (iii) in general it is not idempotent: it may be the case that $\rho_L \circ \rho_L \neq \rho_L$; (iv) the result may depend on the initial order of variables, as shown by the next example.

Example 6 (Incompleteness of the linear explorer). Let us consider an initial state $\bar{s} = \langle \mathbf{v}_0 + \mathbf{v}_1 + \mathbf{v}_2 = 0, \mathbf{v}_3 + \mathbf{v}_1 = 0; \mathbf{v}_2 \in [0, 1], \mathbf{v}_3 \in [0, 1] \rangle$. The reduced state $\rho_{BE}(\bar{s}) = \langle \mathbf{v}_3 + \mathbf{v}_1 = 0, \mathbf{v}_2 + \mathbf{v}_0 - \mathbf{v}_1 = 0; \mathbf{v}_1 \in [-1, 0], \mathbf{v}_2 \in [0, 1], \mathbf{v}_3 \in [0, 1] \rangle$ does not contain the bound $\mathbf{v}_0 \in [-1, 1]$. \square

4.2.2 Combinatorial Explorer (δ_C)

The combinatorial explorer δ_C systematically visits all the bases. It generates all possible combinations of m variables trying to minimize the number of swaps at each basis change. It is very costly, but it finds the best bounds for each variable: it visits all the bases, in particular the one where the optimum is reached. The main advantage with respect to the Simplex is a better tolerance to numerical errors. However it is largely impractical because of (i) the huge cost; and (ii) the negligible gain of precision w.r.t. the use of δ_L that it showed in our benchmark examples.

5 Hints

The relative loss of precision of the join operator incited us to search some ways to recover precision on imprecise operators. We discovered several of those, and they share some properties that led us to define a new concept called hints, which generalizes all of them as well as several known refining techniques.

Hints are precision improving operators which can be used to systematically refine and improve the precision of domain operations in abstract interpretation. Domain operations are either *basic* domain operations (*e.g.*, Υ or \wedge) or their compositions (*e.g.*, $\lambda(\bar{e}_0, \bar{e}_1, \bar{e}_2). (\bar{e}_0 \wedge \bar{e}_1) \Upsilon (\bar{e}_0 \wedge \bar{e}_2)$).

Definition 1 (Hint, h). Let $\diamond \in [C^n \rightarrow C]$ be a concrete domain operation defined over a concrete domain

$\langle C, \sqsubseteq, \sqcup, \sqcap \rangle$. Let $\bar{\diamond} \in [A^n \rightarrow A]$ be the abstract counterpart for \diamond defined over the abstract domain $\langle A, \preceq, \Upsilon, \wedge \rangle$. A hint $\mathfrak{h}_{\bar{\diamond}} \in [A^n \rightarrow A]$ is such that:

$$\begin{aligned} \mathfrak{h}_{\bar{\diamond}}(\bar{e}_0 \dots \bar{e}_{n-1}) &\preceq \bar{\diamond}(\bar{e}_0 \dots \bar{e}_{n-1}) & (\text{Refinement}) \\ \diamond(\gamma(\bar{e}_0) \dots \gamma(\bar{e}_{n-1})) &\sqsubseteq \gamma(\mathfrak{h}_{\bar{\diamond}}(\bar{e}_0 \dots \bar{e}_{n-1})) & (\text{Soundness}). \end{aligned}$$

The first condition states that $\mathfrak{h}_{\bar{\diamond}}$ is a more precise operations than $\bar{\diamond}$. The second condition requires $\mathfrak{h}_{\bar{\diamond}}$ to be a sound approximation of \diamond . An important property of hints is that they can be designed separately and the combined to obtain a more precise hint.

Lemma 1 (Hints combination). *If $\mathfrak{h}_{\bar{\diamond}}^1$ and $\mathfrak{h}_{\bar{\diamond}}^2$ are hints, then*

$$\mathfrak{h}_{\bar{\diamond}}^{\wedge}(\bar{e}_0 \dots \bar{e}_{n-1}) = \mathfrak{h}_{\bar{\diamond}}^1(\bar{e}_0 \dots \bar{e}_{n-1}) \wedge \mathfrak{h}_{\bar{\diamond}}^2(\bar{e}_0 \dots \bar{e}_{n-1})$$

is a hint.

Proof. (Refinement) follows from the definition of \wedge . (Soundness) is because $\diamond(\gamma(\bar{e}_0), \dots, \gamma(\bar{e}_{n-1})) \sqsubseteq \gamma(\mathfrak{h}_{\bar{\diamond}}^1(\bar{e}_0, \dots, \bar{e}_{n-1})) \sqcap \gamma(\mathfrak{h}_{\bar{\diamond}}^2(\bar{e}_0, \dots, \bar{e}_{n-1})) \sqsubseteq \gamma(\mathfrak{h}_{\bar{\diamond}}^{\wedge}(\bar{e}_0, \dots, \bar{e}_{n-1})) \wedge \gamma(\mathfrak{h}_{\bar{\diamond}}^1(\bar{e}_0, \dots, \bar{e}_{n-1})) \wedge \gamma(\mathfrak{h}_{\bar{\diamond}}^2(\bar{e}_0, \dots, \bar{e}_{n-1}))$. \square

The next theorem states that hints improve the precision of static analyses without introducing unsoundness and preserving termination:

Theorem 2 (Refinement of the abstract semantics). *Let \mathfrak{h}_{∇} and \mathfrak{h}_{Υ} be two hints refining respectively the widening and the abstract union, and let \mathfrak{h}_{∇} be a widening operator. Let $\bar{s}^*[\cdot]$ be the abstract semantics obtained from $\bar{s}[\cdot]$ by replacing ∇ with \mathfrak{h}_{∇} and Υ with \mathfrak{h}_{Υ} . Let P be a program. Then, $\forall e \in \mathcal{P}(\Sigma). \forall \bar{e} \in A$.*

$$\begin{aligned} \bar{s}^*[\mathbb{P}](\bar{e}) &\preceq \bar{s}[\mathbb{P}](\bar{e}) & (\text{Refinement}) \\ e \sqsubseteq \gamma(\bar{e}) &\implies \llbracket \mathbb{P} \rrbracket(e) \sqsubseteq \gamma(\bar{s}^*[\mathbb{P}](\bar{e})) & (\text{Soundness}). \end{aligned}$$

Proof sketch. The cases to consider are those for the conditional and the while loop. The conditional can be proven by structural induction. The while loop by instantiating the abstract fixpoint transfer theorem of [12]. \square

5.1 Syntactic hints

Syntactic hints use some part of the program text to refine the operations of the abstract domain. They exploit user annotations to preserve as much information as possible in gathering operations (user-provided hints), and systematically improve the widening heuristics to find tighter loop invariants (thresholds hints).

5.1.1 User-provided hints

They are the easiest, and probably cheapest form of hints. First, we collect all the predicates appearing as assertions or as guards. Then, the gathering operations are refined by explicitly checking for each collected predicate

$$\begin{aligned} \text{pred}(\text{skip};) &= \emptyset \\ \text{pred}(x = E;) &= \emptyset \\ \text{pred}(\text{assert } B;) &= \text{atomize}(B) \\ \text{pred}(\text{assume } B;) &= \text{atomize}(B) \\ \text{pred}(C \ C') &= \text{pred}(C) \cup \text{pred}(C') \\ \text{pred}(\text{if}(B) \{C\} \text{else } \{C'\}) &= \text{atomize}(B) \cup \text{pred}(C) \cup \text{pred}(C') \\ \text{pred}(\text{while}(B) \{C\}) &= \text{atomize}(B) \cup \text{pred}(C) \\ \\ \text{atomize}(B_1 \wedge B_2) &= \text{atomize}(B_1) \cup \text{atomize}(B_2) \\ \text{atomize}(B_1 \vee B_2) &= \text{atomize}(B_1) \cup \text{atomize}(B_2) \\ \text{atomize}(B) &= \{B\} \quad \text{otherwise.} \end{aligned}$$

Fig. 8. The functions `pred` and `atomize` collect the atomic predicates in statements and Boolean expressions.

```
void DomOp()
{ int x = 0, y = 0;
  while (...)
  { if (...) { x++; y += 100; }
    else if (...)
      if (x >= 4) { x++; y++; }
  }
  (* assert x <= y;
  assert y <= 100 * x;
}
```

Fig. 9. Example requiring user-provided hints

B , if it holds for *all* the operands. If this is the case, B is added to the result. The predicate seeker $\text{pred} \in [\text{Stm} \rightarrow \mathcal{P}(\text{BExp})]$ is defined in Fig. 8. User provided hints do not affect the termination of the widening as we can only add finitely many new predicates.

Lemma 2 (User-provided hints).

Let $\diamond \in \{\Upsilon, \nabla\}$, and let P be a program. Then: (i) $\mathfrak{h}_{\diamond}^{\text{pred}}$ defined below is a hint; and (ii) $\mathfrak{h}_{\nabla}^{\text{pred}}$ is a widening operator.

$$\begin{aligned} \mathfrak{h}_{\diamond}^{\text{pred}}(\bar{e}_0, \bar{e}_1) &= \text{let } S = \{B \in \text{pred}(P) \mid \text{A.check}(B, \bar{e}_0) = \text{true} \\ &\quad \wedge \text{A.check}(B, \bar{e}_1) = \text{true}\} \\ &\text{in } \text{A.test}(\bigwedge_{B \in S} B, \diamond(\bar{e}_0, \bar{e}_1)). \end{aligned}$$

Proof sketch. Note that (2.2) implies that $\text{A.test}(b1 \wedge b2, \bar{e}) \preceq \text{A.test}(b1, \bar{e})$, which is enough to prove (Refinement). The soundness condition (2.2) of `check` guarantees that no inconsistent predicate is added to the result, implying (Soundness). \square

Example 7 (Refined SubPoly operations). In example of Fig. 9, $\text{pred}(\text{DomOp}) = \{x \leq y, 4 \leq x, y \leq 100 \cdot x\}$. The refined domain operations keep the predicate $x \leq y$, which is stable among loop iterations, and hence is a loop invariant. \square

We found user-provided hints very useful in `Clousot`, our abstract interpretation based static analyzer for `.Net`. `Clousot` analyzes methods in isolation, and supports assume/guarantee reasoning (“contracts” [30]) via executable annotations [3]. Precision in propagating and

checking program annotations is crucial to provide a satisfactory user experience. User-provided hints help to reach this goal as the analyzer makes sure that at each joint point no user annotation is lost, if it is implied by the incoming abstract states. They make the analyzer more robust w.r.t. incompleteness of Υ or a buggy implementation which may cause Υ to return a more abstract element than the one predicted by the theory. The downside is that user-provided hints are syntactically based:

Example 8 (Fragility of user-provided hints). Let us consider again the code in Fig. 9. If we replace the assertion at (*) with `if 10 <= x then assert 5 <= y`, then $\text{pred}(\text{DomOp}) = \{10 \leq x, 5 \leq y\}$, so that $\text{h}_{\nabla \text{Poly}}^{\text{pred}}$ cannot figure out that $x \leq y$, and hence the analyzer cannot prove that the assertion is valid. Semantic hints (Sect. 5.2.3) will fix it. \square

5.1.2 Thresholds hints

Widening with threshold has been introduced in [5] to improve the precision of standard widenings over non-relational or weakly relational domains. Roughly, the idea of a widening with thresholds is to stage the extrapolation process, so that before projecting a bound to the infinity, values from a set T are considered as candidate bounds. The set T can be either provided by the user or it can be extracted from the program text. The widening with thresholds is just another form of hint. Let \bar{e}_0 and \bar{e}_1 be abstract states belonging to some numerical abstract domain. Without loss of generality we can assume that the basic facts in \bar{e}_0, \bar{e}_1 are in the form $\mathbf{p} \leq k$, where \mathbf{p} is some polynomial. For instance $\mathbf{x} \in [-2, 4]$ is equivalent to $\{-\mathbf{x} \leq 2, \mathbf{x} \leq 4\}$. The standard widening preserves the linear forms with stable upper bounds: $\nabla(\bar{e}_0, \bar{e}_1) = \{\mathbf{p} \leq k \mid \mathbf{p} \leq k_0 \in \bar{e}_0, \mathbf{p} \leq k_1 \in \bar{e}_1, k = \text{if } k_1 > k_0 \text{ then } +\infty \text{ else } k_0\}$. Given a finite set of values T , threshold hints refine the standard widening by:

$$\begin{aligned} \text{h}_{\nabla}^T(\bar{e}_0, \bar{e}_1) = \{ & \mathbf{p} \leq k \mid \mathbf{p} \leq k_0 \in \bar{e}_0, \mathbf{p} \leq k_1 \in \bar{e}_1, \\ & k = \text{if } k_1 > k_0 \text{ then} \\ & \quad \min\{t \in T \cup \{+\infty\} \mid k_1 \leq t\} \\ & \quad \text{else } k_0\}. \end{aligned}$$

Lemma 3. h_{∇}^T is: (i) a hint; and (ii) a widening.

Proof sketch. *Refinement* and *Soundness* are a direct consequence of definition of h_{∇}^T . Termination follows from the fact that T is finite. \square

Example 9 (Widening with thresholds). Let us consider the code snippets in Fig. 10 to be analyzed with Intervals. In the both cases, the (post-)fixpoint is reached after the first iteration $\nabla([0, 0], [1, 1]) = [0, +\infty]$. In the first case, the invariant can be improved by a narrowing step to $\Delta([0, +\infty], [-\infty, 1000]) = [0, 1000]$ (see [10] for

```

void LessThan() {
  int x = 0;
  while (x < 1000) {
    x++;
  }
}
(a) Narrowing

void NotEq() {
  int x = 0;
  while (x != 1000) {
    x++;
  }
}
(b) Thresholds

```

Fig. 10. Two programs to be analyzed with Intervals. The iterations with widening infer the loop invariant $x \in [0, +\infty]$. In the first case, the narrowing step refines the loop invariant to $x \in [0, 1000]$. In the second case, the narrowing fails to refine it.

a definition of narrowing of Intv). In the second case, the narrowing is of no help as $\Delta([0, +\infty], \Upsilon([-\infty, 1000], [1002, +\infty])) = [0, +\infty]$. A widening with Thresholds $T = \{1000\}$ helps discovering the tightest loop invariant for both examples in one step as $\text{h}_{\nabla}^T([0, 0], [1, 1]) = [0, 1000]$. \square

Please note that user-provided hints are of no help in the previous example, as $\text{pred}(\text{NotEq}) = \{x \neq 1000\}$ does not hold for all the operands of the widening.

We are left with problem of generating the set T of thresholds. A common practice in static analyzers is to have $T = \{-1, 0, 1\}$. A better solution is to have the user provide T , left as parameter of the analyzer. This is the approach of [5]. In *Clousot* we chose a slightly different solution, which consists in populating T with the constants appearing in the program text. Constants are fetched from the source using a function $\text{const} \in [\text{Stm} \rightarrow \mathcal{P}(\text{int})]$ defined as one may expect. We found $\text{h}_{\nabla}^{\text{const}}$ very satisfactory. The hint $\text{h}_{\nabla}^{\text{const}}$: (i) helps inferring precise *numerical* loop invariants without requiring the extra iteration steps required for applying the narrowing; and (ii) improves the precision of the analysis of code involving disequalities, *e.g.*, Fig. 10(b). A drawback of threshold hints is that the set T may grow too large, slowing down the convergence of the fixpoint iterations. In *Clousot*, we infer thresholds on a per-method basis, which helps maintaining the cardinality of T quite small.

5.2 Semantic hints

Semantic hints provide a more refined yet more expensive form of operator refinement. For instance, they exploit information in the abstract states to materialize constraints that were implied by the operands (saturation hints, die-hard hints and template hints) or they iterate the application of operators to get a more precise abstract state (reductive hints).

5.2.1 Saturation hints

A common way to design abstract interpreters is to build the abstract domain as a composition of basic abstract domains, which interact through a well-defined interface [13, 6, 20]. Formally, given two abstract domains A_0 ,

A_1 , the Cartesian product $A^\times = A_0 \times A_1$ is still an abstract domain, whose operations are defined as the pointwise extension of those over A_0 and A_1 . Let $\bar{\delta}_i \in [A_i^n \rightarrow A_i]$, $i \in \{0, 1\}$, then

$$\bar{\delta}^\times((\bar{e}_0^0, \bar{e}_1^0) \dots (\bar{e}_0^{n-1}, \bar{e}_1^{n-1})) = (\bar{\delta}_0(\bar{e}_0^0 \dots \bar{e}_0^{n-1}), \bar{\delta}_1(\bar{e}_1^0 \dots \bar{e}_1^{n-1}))$$

The Cartesian product enables the modular design (and refinement) of static analyses. However, a naive design which does not consider the flow of information between the abstract elements may lead to imprecise analyses, as illustrated by the following example.

Example 10 (Cartesian join). Let us consider the abstract domain $Z = \text{Intv} \times \text{LT}$, where $\text{LT} = [\text{Var} \rightarrow \mathcal{P}(\text{Var})]$ is an abstract domain capturing the “less than” relation between *variables*. For instance, $x < y \wedge x < z$ is represented in LT by $[x \mapsto \{y, z\}]$. The domain operations are defined as one may expect [29]. Let $\bar{z}_0 = ([x \mapsto [-\infty, 0], y \mapsto [1, +\infty]], [\cdot])$ and $\bar{z}_1 = ([\cdot], [x \mapsto \{y\}])$ be two elements of Z ($[\cdot]$ denotes the empty map). Then the Cartesian join loses all the information: $\Upsilon^\times(\bar{z}_0, \bar{z}_1) = ([\cdot], [\cdot])$. \square

A common solution is: (i) saturate the operands; and (ii) apply the operation pairwise. The saturation materializes all the constraints implicitly expressed by the product abstract state. Let $\rho \in [A^\times \rightarrow A^\times]$ be a saturation (*a.k.a.* closure) procedure. Then the next lemma provides a systematic way to refine an operator $\bar{\delta}^\times$.

Lemma 4. *The operator $\text{h}_{\bar{\delta}^\times}^\rho$ below is a hint.*

$$\begin{aligned} \text{h}_{\bar{\delta}^\times}^\rho((\bar{e}_0^0, \bar{e}_1^0) \dots (\bar{e}_0^{n-1}, \bar{e}_1^{n-1})) = \\ \text{let } \bar{r}^i = \rho(\bar{e}_0^i, \bar{e}_1^i) \text{ for } i \in 0 \dots n-1 \text{ in } \bar{\delta}^\times(\bar{r}^0 \dots \bar{r}^{n-1}). \end{aligned}$$

Example 11 (Cartesian join, continued). The saturation of \bar{z}_0 materializes the constraint $x < y$: $\bar{r}_0 = ([x \mapsto [-\infty, 0], y \mapsto [1, +\infty], [x \mapsto \{y\}])$, and it leaves \bar{z}_1 unchanged. The constraint $x < y$ is now present in both the operands, and it is retained by the pairwise join. \square

It is worth noting that in general h_∇^ρ does not guarantee the convergence of the iterations, as the saturation procedure may re-introduce constraints which were abstracted away from the widening (*e.g.*, Fig. 10 of [31]).

Saturation hints can provide very precise operations for Cartesian abstract interpretations: They allow the analysis to get additional precision by combining the information present in different abstract domains. The quality of the result depends on the quality of the saturation procedure. The main drawbacks of saturation hints are that: (i) the iteration convergence is not ensured, so that extra care should be put in the design of the widening; (ii) the systematic application of saturation may cause a dramatic slow-down of the analysis. In our experience with the combination of domains implemented in `Clousot`, we found that the slow-down introduced by saturation hints was too high to be practical. Die-hard hints, introduced in the next section, are a better solution to achieve precision without giving up scalability.

5.2.2 Die-hard hints

These hints are based on the observation that often the constraints that one wants to keep at a gathering point often appears explicitly in one of the operands. For instance in Ex. 10 the constraint $x < y$ is explicit in \bar{z}_1 , and implicit in \bar{z}_0 (as $x \leq 0 \wedge 1 \leq y \implies x < y$). Therefore $x < y$ holds for all the operands of the join so it is sound to add it to its result. Die-hard hints generalize and formalize it. They work in three steps: (i) apply the gathering operation, call the result \bar{r} ; (ii) collect the constraints C that are explicit in one of the operands, but are neither present nor implied by \bar{r} ; and (iii) add to \bar{r} all the constraints in C which are implied by *all* the operands. Formally:

$$\begin{aligned} \text{h}_{(\bar{\delta}, I)}^d(\bar{e}_0, \bar{e}_1) &= \text{let } \bar{r} = \bar{\delta}(\bar{e}_0, \bar{e}_1), \\ C &= \cup_{i \in I} \{\kappa \in \bar{e}_i \mid \text{A.check}(\kappa, \bar{r}) = \text{top}\} \\ \text{let } S &= \{\kappa \in C \mid \text{A.check}(\kappa, \bar{e}_0) = \\ &\quad \text{A.check}(\kappa, \bar{e}_1) = \text{true}\} \\ &\text{in } \text{A.test}(\bigwedge_{\kappa \in S} \kappa, \bar{r}). \end{aligned}$$

In defining the die-hard hint for ∇ , one should pay attention to avoid loops which re-introduce a constraint that as been dropped by the widening. One way to do it is to have an asymmetric hint, which restricts C only to the first operand (*e.g.*, the candidate invariant):

Lemma 5. *$\text{h}_{(\Upsilon, \{0,1\})}^d$ and $\text{h}_{(\nabla, \{0\})}^d$ are hints and $\text{h}_{(\nabla, \{0\})}^d$ is a widening.*

5.2.3 Computed hints

Hints can be inferred from the abstract states themselves. By looking at some properties of the elements involved in the operation, one can try to guess useful hints.

Lemma 6 (Computed hints). *Let $\bar{e}_0, \bar{e}_1 \in A$, $\Xi \in [A \times A \rightarrow A]$ a function which returns a set of likely bounds of $\bar{e}_0 \Upsilon \bar{e}_1$. Then h_Υ^Ξ below is a hint.*

$$\begin{aligned} \text{h}_\Upsilon^\Xi(\bar{e}_0, \bar{e}_1) &= \text{let } S = \{\mathbf{B} \in \Xi(\bar{e}_0, \bar{e}_1) \mid \text{A.check}(\mathbf{B}, \bar{e}_0) = \text{true} \\ &\quad \wedge \text{A.check}(\mathbf{B}, \bar{e}_1) = \text{true}\} \\ &\text{in } \text{A.test}(\bigwedge_{\mathbf{B} \in S} \mathbf{B}, \bar{e}_0 \Upsilon \bar{e}_1). \end{aligned}$$

Computed hints are useful when the abstract join Υ is not optimal. Otherwise, h_Υ^Ξ is no more precise than \square . For instance, in a Galois connections-based abstract interpretation, \square is optimal, in that it returns the most precise abstract element overapproximating the concrete union. As a consequence, no further information can be extracted from the operands. It is worth noting that in general h_∇^Ξ is not a widening. However, one can extend the arguments of the previous section to define an asymmetric hint h_∇^Ξ .

The next two kinds of hints (template hints and 2D-convex hull hints) are examples of computed hints.

5.2.4 Template hints

Let $A.\text{range} \in [\text{Exp} \times A \rightarrow \text{Intv}]$ be a function that returns the range for an expression in some abstract state, *e.g.*, it satisfies: $\forall E. \forall \bar{e} \in A. A.\text{range}(E, \bar{e}) = [l, u] \implies \forall \sigma \in \gamma(\bar{e}). l \leq \mathbb{E}[\mathbb{E}](\sigma) \leq u$. If $A.\text{range}(E, \bar{e}_i) = [l_i, u_i]$ for $i \in \{0, 1\}$, then $\gamma(\sqcup_{\text{Intv}}([l_0, u_0], [l_1, u_1]))$ is an upper bound for E in $\cup(\gamma(\bar{e}_0), \gamma(\bar{e}_1))$. As a consequence given a set P of polynomial forms, one can design the guessing function Ξ^P :

$$\begin{aligned} \Xi^P(\bar{e}_0, \bar{e}_1) &= \{l \leq p \leq u \mid p \in P \wedge [l, u] \\ &= \sqcup_{\text{Intv}}(A.\text{range}(p, \bar{e}_0), A.\text{range}(p, \bar{e}_1))\}. \end{aligned}$$

The main difference between $\mathbb{h}_{\gamma}^{\Xi^P}$ and syntactic hints is that the bounds for the polynomials in P are *semantic*, as they are inferred from the abstract states and not from the program text. For instance, computed hints infer the right invariant in Ex. 8 using the set of templates $\text{Oct} \equiv \{x_0 - x_1 \mid x_0, x_1 \text{ are program variables}\}$. In general, template hints with Oct refine SubPoly so to make it as precise as Oct .

5.2.5 2D-Convex Hull hints

New linear inequalities can be discovered at join points using the convex hull algorithm. For instance, the standard join on Poly is defined in that way [14]. However the convex hull algorithm requires an expensive conversion from a tableau of linear constraints to a set of vertices and generators, which causes the analysis time to blow up. A possible solution is to consider a planar convex hull, which computes possible linear relations between *pairs* of variables by: (i) projecting the Intv part of the abstract states on all the two-dimensional planes; and (ii) computing the planar convex hull (of two rectangles, a particularly simple case) on those planes. Planar convex hull, combined with a smart representation of the abstract elements allows us to automatically discover complex invariants without giving up performances.

Example 12 (2D-Convex hull). Let us consider the code in Fig. 11 from [14]. At a price of exponential complexity, Poly can infer the correct loop invariant, and prove the assertion correct. SubPoly refined with 2D-Convex hull hints can prove the assertion, yet keeping a worst-case polynomial complexity [26]. \square

5.2.6 Reductive hints

Intuitively, one way to improve the precision of a unary operator is to iterate its application [18]. However, an unconditional iteration may be source of unsoundness, as shown by the following example.

```
void Foo() {
  int i = 2, j = 0;
  while (...) {
    if (...) { i = i + 4; }
    else     { i = i + 2; j++; } }
  assert 2 <= i - 2 * j; }
```

Fig. 11. Example requiring the use of 2D-convex hull hints to infer the right invariant, expressed by the assertion

Example 13 (Unsoundness of unconditional iterations). Let $- \in [\text{Intv} \rightarrow \text{Intv}]$ be the operator which applies the unary minus to an interval. In general, $\forall n \in \mathbb{N}. \bar{e} = -^{2n}(\bar{e}) \neq -^{2n+1}(\bar{e})$ so that the iterations are unstable. \square

We say that a function f is *reductive* if $\forall x. f(x) \sqsubseteq x$; and *closing* if it is reductive and $\forall x. f(f(x)) = f(x)$.

Lemma 7 (Reductive hints). *Let $\diamond \in [C \rightarrow C]$ be a unary operator and $\bar{\diamond} \in [A \rightarrow A]$ its abstract counterpart. Let \diamond be closing, $\bar{\diamond}$ be reductive, and $n \geq 0$. Then $\mathbb{h}_{\bar{\diamond}}(\bar{e}) = \bar{\diamond}^n(\bar{e})$ is a hint.*

Proof. (Sketch) (Refinement) follows from the definition. To prove *(Soundness)*, it is enough to prove that $\diamond(\gamma(\bar{e})) \subseteq \gamma(\bar{\diamond}^2(\bar{e}))$. It holds as $\diamond(\gamma(\bar{e})) = \bar{\diamond}^2(\gamma(\bar{e})) \subseteq \diamond(\gamma(\bar{\diamond}(\bar{e}))) \subseteq \gamma(\bar{\diamond}^2(\bar{e}))$. \square

The main application of reductive hints is to improve the precision in handling the guards in non-relational abstract domains. Given a Boolean guard B and an abstract domain A , $\psi \equiv \lambda \bar{e}. A.\text{test}(B, \bar{e})$ is an abstract operator which satisfies the hypotheses of Lemma 7. Abstract compilation can be used to express ψ in terms of domain operations, their compositions and state update. Lemma 7 justifies the use of local fixpoint iterations to refine the result of the analysis.

Example 14. Let us consider the following Boolean expression :

$$b1 == b2 \wedge b2 == b3$$

Its abstract compilation in an abstract domain $[\text{Var} \rightarrow \{\text{true}, \text{false}, \top, \perp\}]$ is :

$$\begin{aligned} \psi &\equiv \lambda b. (b[b1, b2 \mapsto b(b1) \wedge b(b2)]) \\ &\quad \hat{\wedge} (b[b2, b3 \mapsto b(b2) \wedge b(b3)]) \end{aligned}$$

where $\hat{\wedge}$ denotes the pointwise extension of \wedge . In an initial abstract state $b_0 = [b1, b2 \mapsto \top; b3 \mapsto \text{true}]$, $\psi(b_0) = [b1 \mapsto \top; b2, b3 \mapsto \text{true}]$, and $\psi^2(b_0) = [b1, b2, b3 \mapsto \text{true}] = \psi^n(b_0)$, $n \geq 2$. \square

6 Refinements for Subpolyhedra

6.1 Precision improvement: Hints

Many of the hints presented above can be used to improve the precision of SubPoly . User-provided hints provide a simple but efficient way to deal with programs

that are not too complicated, and the intervals part of `SubPoly` can of course use the threshold hints.

Saturation hints are impractical for `SubPoly`, but die-hard hints are very useful; indeed, step 3 of the join algorithm (1) can be seen as a particular case of die-hard hints, with the difference that the constraints marked as deleted in the join of the linear equalities domain might not have been actually present in the initial states, but have been introduced in place of an equivalent one by Karr’s algorithm.

Both kinds of computed hints are useful, but in a different way : they can be used for tricky programs that require some complex reasoning, but in most cases the algorithm already returns a good enough result, and use of those hints only slows down the analysis. As an example, template hints can be used to manually set the invariants to infer on a complicated program. They can also be used to guarantee some minimal precision (e.g., at least the precision of octagons) if one can afford the extra time it costs. 2D-convex hull hints are useful in the case of pointer access validation, to infer bounds with non-unary coefficients between pairs of variables.

Reductive hints can also be used on cases involving reduction, because our reduction is not idempotent; however, reduction is rather expensive and there is no guarantee of an actual gain of precision in this case, so we do not use it in practice. Instead, when precision matters, we use reduction with the combinatorial explorer, which does have guarantees on the precision of the result.

6.2 Speed Improvement: Simplification

The simplification operator σ removes redundant information from an abstract element. It is required neither for soundness nor completeness nor to improve the precision of the analysis (unlike ρ), but it is cardinal to the implementation of scalable analyses. The simplification σ of an element of `SubPoly` $\langle \bar{l}; \bar{i} \rangle$ reduces the number of variables in \bar{l} , which is the more expensive domain, without losing any precision. It consists in the application of the following three rules:

- (Const) If an equality $v = b$ is detected, v is projected from \bar{l} and added to \bar{i} ;
- (Slack) If a slack variable β does not appear in $\langle \bar{l}; \bar{i} \rangle$, then it should be removed;
- (Dep) If $\langle \bar{l}; \bar{i} \rangle$ implies $\beta_0 + a \cdot \beta_1 = b$, then one between β_0 and β_1 can be removed.

The rationale behind (Const) is that constants are very expensive when represented in `LinEq` but very cheap if represented with `Intv`; (Slack) performs a kind of garbage collection, by removing slack variables β which are in the domain of $\langle \bar{l}; \bar{i} \rangle$, but such that β does not appear in any of the constraints of \bar{l} and $\bar{i}(\beta) = \top_{\text{Intv}}$; (Dep) is justified by the fact that after refining the intervals for both variables, removing one of the slack variables does

not change the concretization of the abstract element. (Const) is useful when we introduce a new slack variable; (Slack) helps reducing the number of slack variables after joins; and (Dep) is applied as a pre-step of the reduction, to reduce the number of variables and hence make it faster.

7 Experience

We have implemented `SubPoly` on the top of `Clousot`, our modular abstract interpretation-based static analyzer for `.Net`. `Clousot` directly analyzes MSIL, a byte-code target for more than seventy compilers (including C#, Managed C++, VB.NET, F#, ...). Prior to the numerical analysis `Clousot` performs a heap analysis and an expression recovery analysis [28]. `Clousot` performs *intra*-procedural analysis and it supports assume-guarantee reasoning via `Foxtrot` annotations [3, 17]. Contracts are expressed directly in the language as method calls and are persisted to MSIL using the normal compilation process of the source language (cf. Appendix A). Classes and methods are annotated with class invariants, preconditions and postconditions. Preconditions are asserted at call sites and assumed at the method entry point. Postconditions are assumed at call sites and asserted at the method exit point. `Clousot` also checks the absence of specific errors, e.g. out of bounds array accesses, null dereferences, buffer overruns, and divisions by zero.

Figure 12 summarizes our experience in analyzing array creations and accesses in four libraries shipped with `.Net`. The test machine is an ordinary 2.4Ghz dual core machine, running Windows Vista. The assemblies are directly taken from the standard `.NET` directory of our PC. The shipped versions of the assemblies do not contain contracts (We are actively working to annotate the `.Net` libraries). On average, we were able to validate almost 89.5% of the proof obligations. We manually inspected some of the warnings issued for `mscorlib.dll`. Most of them are due to lack of contracts, e.g. an array is accessed using a method parameter or the return value of some helper method. However, we also found real bugs (dead code and off-by-one). That is remarkable considering that `mscorlib.dll` has been tested *in extenso*. We also tried `SubPoly` on the examples of [14, 35, 19], proving all of them.

7.1 Reduction Algorithms

We run the tests using the Simplex-based and the Linear explorer-based reduction algorithms. We used the Simplex implementation shipped with the Microsoft Automatic Graph Layout tool, widely tested and optimized. The results in Fig. 12 show that ρ_{LP} is significantly slower than ρ_{BE} , and in particular the analysis of five

Assembly	Methods	Bounds Checked	Simplex ρ_{LP}			Linear Explorer ρ_{BE}			Max Vars
			Valid	%	Time	Valid	%	Time	
<code>mcorlib.dll</code>	18 084	17 181	14 432	84.00	73:48 (3)	14 466	84.20	23:19 (0)	373
<code>System.dll</code>	13 776	11 891	10 225	85.99	58:15 (2)	10 427	87.69	14:45 (0)	140
<code>System.Web.dll</code>	22 076	14 165	13 068	92.26	24:41 (0)	13 078	92.33	6:33 (0)	182
<code>System.Design.dll</code>	11 419	10 519	10 119	96.20	26:07 (0)	10 148	96.47	5:18 (0)	73
Average				89.00			89.51		

Fig. 12. The experimental results of checking array creation and accesses. `SubPoly` is instantiated with two reductions ρ_{LP} and ρ_{BE} . Time is expressed in minutes, the time-out per method is set to two minutes (in parentheses). The last column reports the maximum number of variables related by an element of `SubPoly`.

methods was aborted as it reached the two minutes time-out. Larger time-outs did not help.

`SubPoly` with the reduction ρ_{LP} validates less accesses than ρ_{BE} . Two reasons for that. First, it is slower, so that the analysis of some methods is aborted and hence some proof obligations cannot be validated. Second, our implementation of the Simplex uses floating point arithmetic which induces some loss of precision. In particular we need to read back the result (a `float`) into an interval of `ints` containing it. In general this may cause a loss of precision and even worse unsoundness. We experienced both of them in our tests. For instance the 39 “missing” proof obligations in `System.Web.dll` and `System.Design.dll` (validated using ρ_{BE} , but not with ρ_{LP}) are due to floating point imprecision in the Simplex. We have considered replacing a floating point-based Simplex with one using exact rationals. However, the Simplex has the tendency to generate coefficients with large denominators. The code we analyze contains many large constants which cause the Simplex to produce enormous denominators.

`SubPoly` with ρ_{BE} instantiated with the linear bases explorer perform very well in practice: it is extremely fast and precise. However, the result may depend on the variables order. A “bad” variable order may cause ρ_{BE} not to infer bounds tight enough. Possible solutions are: (i) to reduce the number of variables using σ (less bases to explore); (ii) to mark variables which can be safely kept in the basis at all times: In the best case, only one basis needs to be explored. In the general case, it still makes the reduction more precise because the bases explored are more likely to give bounds on the variables.

7.2 Max Variables

It is worth noting, that even if `Clousot` performs an intra-procedural analysis, the methods we analyze may be very complex, and they may require tracking linear inequalities among many abstract locations. Abstract locations are produced by the heap analysis [27], and they abstract stack locations and heap locations. Figure 12 shows that it is not uncommon to have methods which requires the abstract state to track more than 100 variables. One single method of `mcorlib.dll` required

to track relations among 373 distinct variables. `SubPoly` handles it: the analysis with ρ_{BE} took a little bit more than a minute. To the best of our knowledge those performances in presence of so many variables are largely beyond state-of-the-art `Poly` implementations.

7.3 Hints

Figure 13 focuses on the analysis of `mcorlib` using `SubPoly` refined with hints and no-reduction. The first column in the table shows the results of the analysis with no hints. This is roughly equivalent to precisely propagating arbitrary linear equalities and intervals, with limited inference and no propagation of information between linear equalities and intervals. User-provided hints and die-hard hints add more inference power, at the price of a still acceptable slow-down. Computed hints (with Octagons and 2D-Convex hull) further slow-down of the analysis, causing the analysis of various methods to time out. We manually inspected the analysis logs to investigate the differences. Ignoring the methods that timed-out, with respect to `SubPoly*`, $\langle \text{SubPoly}^*, \mathbb{h}_{\bar{v}}^{\text{Oct}} \rangle$ and $\langle \text{SubPoly}^*, \mathbb{h}_{\bar{v}}^{\text{2DCH}} \rangle$ report respectively 125 and 124 less false positives. Out of those, only 13 overlap.

One may wonder if computed hints are needed at all. We observed that, when considering annotated code (unfortunately, just a small fraction of the overall codebase at the moment of writing), one needs to refine the operations of the abstract domains with hints in order to get a very low (and hence acceptable) false alarms ratio (around 0.5%). In fact, even if (relatively) rare, assertions as in Fig. 7(b) and Fig. 11 are present in real code. Thanks to the incremental structure of `Clousot`, we do not need to run `SubPoly` with *all* the hints on *all* the analyzed methods, but we can focus the highest precision only on the few methods which require it.

8 Conclusions

We introduced `SubPoly`, a new numerical abstract domain based on the combination of linear equalities and intervals. `SubPoly` can track linear inequalities involving hundred of variables. We defined the operations of

SubPoly		SubPoly*		Slow	SubPoly* + $\mathfrak{h}_{\overline{\gamma}}^{\text{Oct}}$		Slow	SubPoly* + $\mathfrak{h}_{\overline{\gamma}}^{\text{2DCH}}$		Slow
Valid	Time	Valid	Time	down	Valid	Time	down	Valid	Time	down
14 230	4:29(0)	14 432	20:22(0)	4.5x	13 948	81:24(20)	18.2x	14 396	36:33(7)	8.1x

Fig. 13. The experimental results analyzing `mscorlib` with `SubPoly` and different semantic hints and no-reduction. `SubPoly*` denotes `SubPoly` refined with $\mathfrak{h}_{\delta}^{\text{pred}}$ and $\mathfrak{h}_{\overline{\gamma}, \gamma}^d$. Computed hints significantly slow-down the analysis, but they are needed to reach a very low false alarm ratio.

the abstract domain (order, join, meet, widening); the simplification operator (to speed up the analysis); and two reduction operators (one based on linear programming and another based on basis exploration). We found Simplex-based reduction quite unsatisfactory for program analysis purposes: because of floating point errors the result may be too imprecise, or worse unsound. We introduced then the basis exploration-based reduction, in practice more precise and faster.

`SubPoly` precisely propagates linear inequalities, but it may fail to infer some of them at join points. Precision can be recovered using hints either provided by the programmer in the form of program annotations; or automatically generated (at some extra cost). `SubPoly` worked fine on some well known examples in literature that required the use of `Poly`. We tried `SubPoly` on shipped code, and we showed that it scales to several hundreds of variables, a result far beyond existing `Poly` implementations.

Acknowledgments Thanks to Lev Nachmanson for providing us the Simplex implementation. Thanks to Manuel Fähndrich, Jérôme Feret, Corneliu Popeea for the useful discussions.

References

- R. Bagnara, P.M. Hill, and E. Zaffanella. The Parma Polyhedra Library. <http://www.cs.unipr.it/pp1/>.
- M. Barnett, M. Fähndrich, and F. Logozzo. Embedded contract languages. In *SAC'10*. ACM Press, 2010.
- M. Barnett, M. A. Fähndrich, and F. Logozzo. Fox-trot and Clousot: Language agnostic dynamic and static contract checking for .Net. Technical Report MSR-TR-2008-105, Microsoft Research, 2008.
- M. Barnett, M. Fähndrich, D. Garbervetsky, and F. Logozzo. Annotations for (more) precise points-to analysis. In *IWACO 2007*, jul 2007.
- 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 *PLDI'03*. ACM Press.
- B.-Y. E. Chang and K. R. M. Leino. Abstract interpretation with alien expressions and heap structures. In *VMCAI'05*. Springer-Verlag, 2005.
- V. Chvátal. *Linear Programming*. W. H. Freeman, 1983.
- R. Clarisó and J. Cortadella. The octahedron abstract domain. In *SAS'04*.
- P. Cousot. The calculational design of a generic abstract interpreter. In *Computational System Design*. NATO ASI Series F. IOS Press, Amsterdam, 1999.
- P. Cousot and R. Cousot. Abstract interpretation: a unified lattice model for static analysis of programs by construction or approximation of fixpoints. In *POPL'77*.
- P. Cousot and R. Cousot. Systematic design of program analysis frameworks. In *POPL '79*.
- P. Cousot and R. Cousot. Abstract interpretation and application to logic programs. *Journal of Logic Programming*, 13(2–3):103–179, July 1992.
- P. Cousot, R. Cousot, J. Feret, L. Mauborgne, A. Miné, D. Monniaux, and X. Rival. Combination of abstractions in the ASTRÉE static analyzer. In *ASIAN'06*, volume 4435 of *LNCS*, pages 272–300. Springer-Verlag.
- P. Cousot and N. Halbwegs. Automatic discovery of linear restraints among variables of a program. In *POPL '78*.
- G. B. Dantzig. Programming in linear structures. Technical report, USAF, 1948.
- M. Fähndrich and F. Logozzo. Static contract checking with abstract interpretation. In *FoVeOOS'10*, LNCS. Springer-Verlag, 2010.
- P. Ferrara, F. Logozzo, and M. A. Fähndrich. Safer unsafe code in .Net. In *OOPSLA'08*.
- P. Granger. Improving the results of static analyses programs by local decreasing iteration. In *FSTTCS'92*. Springer-Verlag.
- B. S. Gulavani, S. Chakraborty, A. V. Nori, and S. K. Rajamani. Automatically refining abstract interpretations. In *TACAS'08*.
- S. Gulwani, B. McCloskey, and A. Tiwari. Lifting abstract interpreters to quantified logical domains. In *POPL'08*. ACM Press, 2008.
- S. Gulwani, S. Srivastava, and R. Venkatesan. Program analysis as constraint solving. In *PLDI'08*.
- M. Karr. On affine relationships among variables of a program. *Acta Informatica*, 6(2):133–151, July 1976.
- L. Khachiyan, E. Boros, K. Borys, K. M. Elbassioni, and M. Gurvich. Generating all vertices of a polyhedron is hard. In *SODA'06*.
- L. Kovács. Reasoning algebraically about p-solvable loops. In *TACAS'08*. Springer-Verlag, 2008.
- V. Laviro and F. Logozzo. Refining abstract interpretation-based static analyses with hints. In *APLAS'09*.
- V. Laviro and F. Logozzo. Subpolyhedra: a (more) scalable approach to infer linear inequalities. In *VMCAI'09*.
- F. Logozzo. Cibai: An abstract interpretation-based static analyzer for modular analysis and verification of Java classes. In *VMCAI'07*, 2007.
- F. Logozzo and M. A. Fähndrich. On the relative completeness of bytecode analysis versus source code analysis. In *CC'08*.
- F. Logozzo and M. A. Fähndrich. Pentagons: A weakly relational abstract domain for the efficient validation of array accesses. In *SAC'08*.

30. B. Meyer. *Object-Oriented Software Construction (2nd Edition)*. Professional Technical Reference. Prentice Hall, 1997.
31. A. Miné. The octagon abstract domain. In *WCRE 2001*.
32. A. Miné. *Weakly Relational Numerical Abstract Domains*. PhD thesis, École Polytechnique, 2004.
33. M. Müller-Olm and H. Seidl. Precise interprocedural analysis through linear algebra. In *POPL'04*.
34. E. Rodríguez-Carbonell and D. Kapur. Automatic generation of polynomial invariants of bounded degree using abstract interpretation. *Sci. Comput. Program.*, 64(1), 2007.
35. S. Sankaranarayanan, F. Ivancic, and A. Gupta. Program analysis using symbolic ranges. In *SAS'07*.
36. S. Sankaranarayanan, H. B. Sipma, and Z. Manna. Scalable analysis of linear systems using mathematical programming. In *VMCAI'05*.
37. A. Simon. *Value-Range analysis of C Programs*. Springer-Verlag, 2008.
38. A. Simon, A. King, and J. Howe. Two variables per linear inequality as an abstract domain. In *LOPSTR'02*.
39. D. A. Spielman and S.-H. Teng. Smoothed analysis of algorithms: Why the simplex algorithm usually takes polynomial time. *J. ACM.*, 51(3), 2004.

A Foxtrot

Foxtrot is a language independent solution for contract specifications in `.Net`. It does not require any source language support or compiler modification. Preconditions and postconditions are expressed by invocations of static methods (`Contract.Requires` and `Contract.Ensures`) at the start of methods. Class invariants are contained in a method with an opportune name (`ObjectInvariant`) or tagged by a special attribute (`[ObjectInvariant]`). Dummy static methods are used to express meta-variables such as *e.g.* `Contract.Old(x)` for the value in the pre-state of `x` or `Contract.WritableBytes(p)` for the length of the memory region associated with `p`. These contracts are persisted to MSIL using the standard source language compiler.

Contracts in the **Foxtrot** notation (using static method calls) can express arbitrary boolean expressions as preconditions and post-conditions. We expect the expressions to be side effect free (and only call side-effect free methods). We use a separate purity checker to optionally enforce this [4].

A binary rewriter tool enables dynamic checking. It extracts the specifications and instruments the binary with the appropriate runtime checks at the applicable program points, taking contract inheritance into account. Most **Foxtrot** contracts can be enforced at runtime.

For static checking, **Foxtrot** contracts are presented to **Clousot** as simple `assert` or `assume` statements. *E.g.*, a pre-condition of a method appears as an assumption at the method entry, whereas it appears as an assertion at every call-site.

B Simplex algorithm

We recall some basic facts about the Simplex algorithm, and in particular the notion of basis. The Simplex algorithm finds the best solution to the problem:

$$\begin{aligned} & \text{maximize} && c^T \mathbf{v} \\ & \text{subject to} && A \mathbf{v} = \mathbf{b} \end{aligned}$$

There may be also interval constraints ($l_i \leq v_i \leq u_i$), but they are not important for the notion of basis. The problem above can be rewritten in matricial form as

$$(A \mid \mathbf{b}) \begin{pmatrix} \mathbf{v} \\ -1 \end{pmatrix} = 0$$

We let $S = (A|\mathbf{b})$. There are infinitely many matrices S with the same space of solutions as $S \mathbf{v} = 0$, so we can make a few assumptions on S . First, we can use Gaussian elimination get an upper triangular matrix (row echelon form). Gaussian elimination updates the matrix by adding to a row a linear combination of the other rows of the matrix, which does not change the space of solutions; after several of such updates, the result is triangular. We can then remove all zero rows and divide each row by its leading coefficient (which is the left-most non-zero coefficient). These operations do not change the space of solutions. As Gaussian elimination guarantees that the leading coefficient of each row is strictly right of the leading coefficients of the rows above it, there is at most one leading coefficient in each column. The variables whose columns contain a leading coefficient are called *basic variables*, the ones whose columns do not contain a leading coefficient are called *non-basic variables*. The set of basic variables is the *basis*. It is also convenient to have the columns corresponding to basic variables containing only zeros except for a single one (the leading coefficient). This can be achieved from the previous matrix by a way similar to Gaussian elimination.

The Simplex algorithm starts with a matrix in this form, and at each iteration changes the basis. *Changing the basis* consists in choosing a basic variable, v_b , with the associated row r (the row whose leading coefficient is in the column for v_b), then choosing a non-basic variable v_n whose coefficient c in the row r is non-zero, then divide the row r by c , and use row operations to make all the other coefficients in the column for v_n zeros. The basis is now the previous basis plus v_n minus v_b (and so v_b is now non-basic and v_n is now basic). Note that the matrix may not be triangular anymore; this is not required for the simplex algorithm. The simplex algorithm uses the cost function c and the bounds l_i and u_i on variables to change the basis. Furthermore the simplex chooses the variables to ensures that c will not be zero; if this is not the case, a zero coefficient means that a particular exchange is not possible.