

Refining Abstract Interpretation-based Static Analyses with Hints

Vincent Laviron¹ Francesco Logozzo²

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

`Vincent.Laviron@ens.fr`

² Microsoft Research, Redmond, WA (USA)

`logozzo@microsoft.com`

Abstract. We focus our attention on the loss of precision induced by abstract domain *operations*. We introduce a new technique, *hints*, which allows us to systematically refine the operations defined over elements of an abstract domain. We formally define hints in the abstract interpretation theory, we prove their soundness, and we characterize two families of hints: syntactic and semantic. We give some examples of hints, and we provide our experience with hints in `Clousot`, our abstract interpretation-based static analyzer for `.Net`.

1 Introduction

The three main elements of an abstract interpretation are: (i) the abstract elements (“*which properties am I interested in?*”); (ii) the abstract transfer functions (“*which is the abstract semantics of basic statements?*”); and (iii) the abstract operations (“*how do I combine the abstract elements?*”).

The loss of precision induced by the abstract elements is exemplified by Fig. 1(a). The assertion cannot be proved using only convex numerical abstract domains such as Intervals [4], Pentagons [18], Octagons [23] or even Polyhedra [8]. The reason for that is that the most precise property at the join point $x == -1 \vee x == 1$ cannot be exactly represented in any of those domains. For instance Intervals (`Intv`) approximate it with $-1 \leq x \leq 1$, so that the fact that $x \neq 0$ is lost. Many techniques have been proposed to overcome this problem. They essentially rely on the refinement of the *elements* of the abstract domain. Solutions include trace partitioning [16, 20, 9], domain completion [5], powerset construction [1, 19], and abstract domain extension [25]. Abstract transfer functions may introduce an orthogonal loss of precision. For instance, in Fig. 1(b) the expression initializer for `z` is in a quadratic form. Thus no linear numerical abstract domain can precisely capture the relation between `x`, `y` and `z`. Standard domain refinements are of no help. A rough transfer function can simply abstract away `z`. A more precise one can approximate `z` with an interval. However, a compositional evaluation of the expressions which mimics the concrete one is not precise enough to discharge the assertion $-2 \leq z$. Several authors suggested methods to infer optimal transfer functions in particular settings as, *e.g.*, constraint matrices [24], shape analysis [26] or constant propagation [3].

```

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)

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;
}
(c)

```

Fig. 1. Examples of orthogonal losses of precision in abstract interpretations: (a) a convex domain cannot represent $x \neq 0$; (b) a compositional transfer function does not infer the tightest lower bound for z ; and (c) the standard domain operations on Polyhedra are not precise enough to infer the loop invariant $x \leq y$.

Surprisingly enough, the refinement of the *operations* over abstract elements has been widely ignored in the literature (with the exceptions of [14, 1, 11, 12] which however focused their attention just on the widening operator).

Example Let us consider the code snippet in Fig. 1(c). In order to prove the assertions valid, the static analysis should infer the loop invariant $x \leq y \leq 100 \cdot x$. Different abstract domains infer different invariants. The Octagon abstract domain (**Oct**) is a *weakly relational* domain which captures properties in the form of $\pm x \pm y \leq k$. It infers the loop invariant: $0 \leq x \wedge 0 \leq y \wedge x \leq y$. The Polyhedra abstract domain (**Poly**) is a *fully relational* domain. It can infer and represent arbitrary linear inequalities: Abstract elements are in the form $\sum_i a_i \cdot x_i \leq k$. As a consequence one expects **Poly** to always be more precise than **Oct**. However, when applied to the example, **Poly** infers the loop invariant: $0 \leq x \wedge 0 \leq y \wedge y \leq 100 \cdot x$: Even if **Poly** can *exactly* represent the constraint $x \leq y$, it fails inferring it! This is quite surprising. The reason for that should be found in the widening operators over the two domains. The (standard) widening over Octagons explicitly seeks an upper bound for the difference $x - y$ (in the example 0). The (standard) widening over Polyhedra preserves the inequalities that are stable over two loop iterations. In the example, the constraint $x \leq y$, even if implied by the abstract states to be widened, is never materialized. Therefore, the state after widening does not contain it either. Intuitively, in order to obtain the most precise loop invariant for **DomOp**, one needs to refine the widening operator for **Poly** to be at least as precise as the one for **Oct**. One way to refine the widening is by remarking that the predicate $x \leq y$ appears as a condition of some assertion, and then trying to explicitly materialize it. Another possible refinement is by seeking upper bounds for the expression $x - y$. The first is an example of *syntactic hint*. The latter is an example of *semantic hint*. Those observations can be extended and generalized to other abstract domain operators.

The case for hints The main goal for a static analysis designer is the precision/speed trade-off. To achieve it, the common practice is to drop some of

the expressive power of the analysis while maximizing the inference power. In `Clousot`, our static analyzer for `.Net`, we needed additional flexibility. `Clousot` is mainly used to validate code contracts expressed by users in form of pre-conditions, post-conditions and object invariant. First, we observed that usual weakly-relational abstract domains are not precise enough to be used in a modular checker: for instance, it is often the case that the argument to establish an “easy” precondition (*e.g.*, $x \leq y$) at the call site involves a complex reasoning between several linear inequalities which require expensive abstract domains as `Poly`. Second, we needed `Clousot` to be adaptable, in that it can either run in an interactive environment (faster, but with more noise) or on a build machine overnight (slower, but much more precise). As a consequence, we took a different direction in the design of the abstract domains in `Clousot`: we retained the *expressive* power while we gave up some of the *inference* (*e.g.*, Pentagons [18] and Subpolyhedra [17]). Hints, introduced in this paper, are an orthogonal to the abstract domain, and they allow us to incrementally increase the precision of the analysis by refining the transfer functions.

Contribution We introduce hints, a new technique which allows us to systematically refine static analyses. The main ideas of hints are: (i) to have a separate module to figure out which constraints or families of constraints are of interest for the analysis; and (ii) to use such a module to refine the *operations* of the abstract domain. The main difference with related works on automatic refinement of static analyses is that hints refine the *operations* over abstract elements and *not* the elements themselves nor the transition relations. The main advantages of hints are that: (i) they enable an easy refinement of static analyses; (ii) they enable a fine-tuning of the cost/performance ratio; (iii) they make the analysis more robust with respect to implementation-related precision bugs. Hints are useful when the abstract operations are not complete w.r.t. the concrete ones, which is often the case in practice.

We formalize hints using the abstract interpretation theory, and we prove them correct w.r.t. a generic abstract interpreter. We characterize syntactic (user-defined, thresholds) and semantic hints (saturation, die-hard, computed, reductive). We show how they generalize existing techniques as, *e.g.*, widening with thresholds [2]. We applied hints to SubPolyhedra (`SubPoly`), a new, very efficient numerical abstract domain to propagate arbitrary linear inequalities. `SubPoly` has the same expressive power as `Poly`, but drops some of the inference to achieve scalability. Hints allow `SubPoly` to recover precision without giving up performances. Hints are implemented in `Clousot`, our static analyzer for `.NET` available at [21].

2 Abstract Interpretation Frameworks

Abstract interpretation is a general theory of approximations which formalizes the intuition that the semantics of a program is more or less precise depending on the observation level. In particular, the static analysis of a program is a semantics precise enough to capture the properties of interest and coarse enough

to be computable. The *concrete* and *abstract* semantics of a program are defined as fixpoints respectively over a concrete and an abstract domain. The concrete and the abstract domains are related by a soundness relation, which induces the soundness of the abstract semantics [6].

Static approximation: Abstract Domain In the Galois connections approach to abstract interpretation [4], the concrete domain and the abstract domain are assumed to be two *complete* lattices, respectively $\langle C, \sqsubseteq, \sqcup \rangle$ and $\langle A, \sqsubseteq, \sqcup \rangle$. The soundness relation is expressed by a pair of monotonic functions $\langle \alpha, \gamma \rangle$, such that $\forall e \in C. \forall \bar{e} \in A. \alpha(e) \sqsubseteq \bar{e} \iff e \sqsubseteq \gamma(\bar{e})$. In such a setting, the abstract join operator \sqcup is optimal in that: $\forall e_1, e_2 \in C. \alpha(\sqcup(e_1, e_2)) = \sqcup(\alpha(e_1), \alpha(e_2))$ [5]. In practice, most analyses do not require the existence of the *best* approximation for concrete elements, a sound approximation suffices. For instance, there is no best polyhedron approximating the set of concrete points $B = \{(x, y) \in \mathbb{R}^2 \mid x^2 + y^2 \leq 1\}$. However, any polyhedron including B is a sound abstraction. In the relaxed form of abstract interpretation [6], the abstract domain is not required to be complete under \sqcup . It is simply a *pre-order* $\langle A, \preceq, \Upsilon, \wedge \rangle$. The soundness relation is expressed by a monotonic concretization function $\gamma \in [A \rightarrow C]$, *i.e.*, no abstraction function is required. The abstract union Υ gathers together the information flowing from incoming edges. It is not required to be the *least* upper bound (which may not exist at all): $\forall \bar{e}_0, \bar{e}_1 \in A. \bar{e}_0 \preceq \bar{e}_0 \Upsilon \bar{e}_1 \wedge \bar{e}_1 \preceq \bar{e}_0 \Upsilon \bar{e}_1$. It is a sound, but not *optimal*, approximation of the concrete join: $\forall \bar{e}_0, \bar{e}_1 \in A. \sqcup(\gamma(\bar{e}_0), \gamma(\bar{e}_1)) \sqsubseteq \gamma(\Upsilon(\bar{e}_0, \bar{e}_1))$. The abstract intersection returns a common lower bound for the operands, which approximates the concrete meet: $\forall \bar{e}_0, \bar{e}_1 \in A. \gamma(\bar{e}_0) \sqcap \gamma(\bar{e}_1) \sqsubseteq \gamma(\bar{e}_0 \wedge \bar{e}_1)$.

Hereafter we assume: (i) the concrete domain to be the complete lattice $\langle \mathcal{P}(\Sigma), \subseteq, \cup, \cap \rangle$ where Σ is a set of concrete program states mapping variables to values; (ii) the abstract domain to be a pre-order $\langle A, \preceq, \Upsilon, \wedge \rangle$, therefore putting ourselves in the setting of the relaxed form of abstract interpretation.

Dynamic approximation: Widening/Narrowing In general A is of infinite height, so that the fixpoint computation may not terminate. A widening operator ∇ should then be defined to ensure the convergence of the iteration to a *post*-fixpoint. Formally ∇ satisfies: (i) $\forall \bar{e}_0, \bar{e}_1 \in A. \bar{e}_0, \bar{e}_1 \preceq \nabla(\bar{e}_0, \bar{e}_1)$; and (ii) for each (possibly infinite) sequence of abstract elements $\bar{e}_0, \bar{e}_1 \dots \bar{e}_k$ the sequence defined by $\bar{e}_0^\nabla = \bar{e}_0, \bar{e}_1^\nabla = \nabla(\bar{e}_0^\nabla, \bar{e}_1) \dots \bar{e}_k^\nabla = \nabla(\bar{e}_{k-1}^\nabla, \bar{e}_k)$ is ultimately stationary. It is worth noting that a widening operator is not commutative. The loss of precision introduced by the widening can be partially recovered using a narrowing operator. A narrowing Δ operator satisfies: (i) $\forall \bar{e}_0, \bar{e}_1 \in A. \wedge(\bar{e}_0, \bar{e}_1) \preceq \Delta(\bar{e}_0, \bar{e}_1) \preceq \bar{e}_0, \bar{e}_1$; and (ii) for each (possibly infinite) sequence of abstract elements $\bar{e}_0, \bar{e}_1 \dots \bar{e}_k$ the sequence defined by $\bar{e}_0^\Delta = \bar{e}_0, \bar{e}_1^\Delta = \Delta(\bar{e}_0^\Delta, \bar{e}_1) \dots \bar{e}_k^\Delta = \Delta(\bar{e}_{k-1}^\Delta, \bar{e}_k)$ is ultimately stationary.

Transfer functions It is common practice for the implementation of an abstract domain A to provide some primitive transfer functions. The assignment abstract transfer function, $A.\text{assign}$, 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 : $\forall x, E. \forall \bar{e} \in A. \{\sigma[x \mapsto v] \mid \sigma \in \gamma(\bar{e}), \mathbb{E}[\mathbb{E}](\sigma) = v\} \subseteq$

$$\begin{aligned}
\bar{s}[\text{skip};](\bar{e}) &= \bar{e} & \bar{s}[\mathbf{x} = \mathbf{E};](\bar{e}) &= \mathbf{A}.\text{assign}(\bar{e}, \mathbf{x}, \mathbf{E}) \\
\bar{s}[\text{assume } \mathbf{B};](\bar{e}) &= \bar{s}[\text{assert } \mathbf{B};](\bar{e}) = \mathbf{A}.\text{test}(\mathbf{B}, \bar{e}) \\
\bar{s}[\mathbf{C}_1 \ \mathbf{C}_2](\bar{e}) &= \bar{s}[\mathbf{C}_2](\bar{s}[\mathbf{C}_1](\bar{e})) \\
\bar{s}[\text{if } (\mathbf{B}) \ \{\mathbf{C}_1\} \text{else } \{\mathbf{C}_2\};](\bar{e}) &= \bar{s}[\mathbf{C}_1](\mathbf{A}.\text{test}(\mathbf{B}, \bar{e})) \vee \bar{s}[\mathbf{C}_2](\mathbf{A}.\text{test}(!\mathbf{B}, \bar{e})) \\
\bar{s}[\text{while}(\mathbf{B}) \ \{\mathbf{C}\};](\bar{e}) &= \text{let } \bar{I} = \text{fix } \lambda X. \bar{e} \vee \bar{s}[\mathbf{C}](\mathbf{A}.\text{test}(\mathbf{B}, X)) \\
&\quad \text{in } \mathbf{A}.\text{test}(!\mathbf{B}, \bar{I})
\end{aligned}$$

Fig. 2. The abstract semantics for the while-language

$$\begin{aligned}
\llbracket \text{skip}; \rrbracket(\mathbf{e}) &= \mathbf{e} & \llbracket \mathbf{x} = \mathbf{E}; \rrbracket(\mathbf{e}) &= \{\sigma[\mathbf{x} \mapsto v] \mid \sigma \in \mathbf{e}, \mathbb{E}[\mathbf{E}](\sigma) = v\} \\
\llbracket \text{assume } \mathbf{B}; \rrbracket(\mathbf{e}) &= \llbracket \text{assert } \mathbf{B}; \rrbracket(\mathbf{e}) = \{\sigma \in \mathbf{e} \mid \mathbb{B}[\mathbf{B}](\sigma) = \text{true}\} \\
\llbracket \mathbf{C}_1 \ \mathbf{C}_2 \rrbracket(\mathbf{e}) &= \llbracket \mathbf{C}_2 \rrbracket(\llbracket \mathbf{C}_1 \rrbracket(\mathbf{e})) \\
\llbracket \text{if } (\mathbf{B}) \ \{\mathbf{C}_1\} \text{else } \{\mathbf{C}_2\}; \rrbracket(\mathbf{e}) &= \llbracket \mathbf{C}_1 \rrbracket(\{\sigma \in \mathbf{e} \mid \mathbb{B}[\mathbf{B}](\sigma) = \text{true}\}) \cup \llbracket \mathbf{C}_2 \rrbracket(\{\sigma \in \mathbf{e} \mid !\mathbb{B}[\mathbf{B}](\sigma) = \text{true}\}) \\
\llbracket \text{while}(\mathbf{B}) \ \{\mathbf{C}\}; \rrbracket(\bar{\mathbf{e}}) &= \text{let } \mathbf{I} = \bigcup_i \llbracket \mathbf{C} \rrbracket^i(\{\sigma \in \mathbf{I} \mid \mathbb{B}[\mathbf{B}](\sigma) = \text{true}\}) \\
&\quad \text{in } \{\sigma \in \mathbf{I} \mid !\mathbb{B}[\mathbf{B}](\sigma) = \text{true}\}.
\end{aligned}$$

Fig. 3. The reachable states semantics for the while-language.

$\gamma(\mathbf{A}.\text{assign}(\bar{\mathbf{e}}, \mathbf{x}, \mathbf{E}))$. The test abstract transfer function, $\mathbf{A}.\text{test}$, filters the input states ($\mathbb{B}[\mathbf{B}](\sigma)$ denotes the evaluation of the Boolean expression \mathbf{B} in the state σ): $\forall \mathbf{B}.\forall \bar{\mathbf{e}} \in \mathbf{A}.\{\sigma \in \gamma(\bar{\mathbf{e}}) \mid \mathbb{B}[\mathbf{B}](\sigma) = \text{true}\} \subseteq \gamma(\mathbf{A}.\text{test}(\bar{\mathbf{e}}, \mathbf{B}))$. The abstract checking $\mathbf{A}.\text{check}$ verifies if an assertion \mathbf{A} holds in an abstract state $\bar{\mathbf{e}}$. It has four possible outcomes: *true* (\mathbf{A} holds in all the concrete states $\gamma(\bar{\mathbf{e}})$); *false* (\mathbf{A} holds in all the concrete states $\gamma(\bar{\mathbf{e}})$); *bottom* (the assertion is unreachable); *top* (the validity of \mathbf{A} cannot be decided in $\gamma(\bar{\mathbf{e}})$).

3 Concrete and Abstract semantics for a While language

We illustrate hints on a simple abstract interpreter for a while language. The concrete, reachable states semantics $\llbracket \cdot \rrbracket \in [\mathbf{Stm} \times \mathcal{P}(\Sigma) \rightarrow \mathcal{P}(\Sigma)]$ is in Fig. 3. The abstract semantics $\bar{s}[\cdot] \in [\mathbf{Stm} \times \mathbf{A} \rightarrow \mathbf{A}]$ is in Fig. 2. It is parametrized by the abstract domain $\langle \mathbf{A}, \preceq, \vee, \wedge \rangle$ and a set of primitives `assign`, `test`, and `check`. The `skip` statement has no effect on the abstract state. The effects of the assignment, the assumption and the assertion are handled by the corresponding primitives of the abstract domain. Please note that for the purposes of the analysis the effects of `assume` and `assert` coincide: The assertions will be checked in a second phase, after the analysis has inferred the program invariants for all the program points. The abstract semantics of sequence is function composition. The abstract semantics of conditional: (i) pushes the guard and its negation onto the two branches; and (ii) gathers the effects using the abstract union. The abstract semantics of `while` is given in terms of `fix`, which computes the loop invariant as the limit of the fixpoint iterations with widening. Given a function $F \in [\mathbf{A} \rightarrow \mathbf{A}]$, $\text{fix}(F)$ is the limit of the iteration sequence: $I^0 = \perp$; $I^{n+1} = \text{if } F(I^n) \preceq I^n \text{ then } I^n \text{ else } I^n \vee F(I^n)$. The post-state for `while` is then obtained by intersecting the loop invariant

with the negation of the guard. It is easy to show that for any program P , $\forall e \in \mathcal{P}(\Sigma). \forall \bar{e} \in \mathbf{A}. e \subseteq \gamma(\bar{e}) \implies \llbracket P \rrbracket(e) \subseteq \gamma(\bar{s}\llbracket P \rrbracket(\bar{e}))$.

4 Hints

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, \mathfrak{h}). Let $\diamond \in [\mathbf{C}^n \rightarrow \mathbf{C}]$ be a concrete domain operation defined over a concrete domain $\langle \mathbf{C}, \sqsubseteq, \sqcup, \sqcap, \sqsupset \rangle$. Let $\bar{\diamond} \in [\mathbf{A}^n \rightarrow \mathbf{A}]$ be the abstract counterpart for \diamond defined over the abstract domain $\langle \mathbf{A}, \preceq, \Upsilon, \wedge \rangle$. A hint $\mathfrak{h}_{\bar{\diamond}} \in [\mathbf{A}^n \rightarrow \mathbf{A}]$ 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. Therefore, if $\mathfrak{h}_{\bar{\diamond}_1}^1$ and $\mathfrak{h}_{\bar{\diamond}_2}^2$ are hints, then $\mathfrak{h}_{\bar{\diamond}}^{\wedge}(\bar{e}_0 \dots \bar{e}_{n-1}) = \mathfrak{h}_{\bar{\diamond}_1}^1(\bar{e}_0 \dots \bar{e}_{n-1}) \wedge \mathfrak{h}_{\bar{\diamond}_2}^2(\bar{e}_0 \dots \bar{e}_{n-1})$ is a hint, too. Hints improve the precision of static analyses without introducing unsoundness and preserving termination:

Theorem 1 (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, $\forall e \in \mathcal{P}(\Sigma). \forall \bar{e} \in \mathbf{A}$.

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

5 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).

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

<pre> void LessThan() { int x = 0; while (x < 1000) x++; } </pre> <p style="text-align: center;">(a) Narrowing</p>	<pre> void NotEq() { int x = 0; while (x != 1000) x++; } </pre> <p style="text-align: center;">(b) Thresholds</p>
---	---

Fig. 4. 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.

Lemma 1 (User-provided hints). *Let $\diamond \in \{\Upsilon, \nabla\}$, and let P be a program. Then: (i) $\mathbb{h}_\diamond^{\text{pred}}$ defined below is a hint; and (ii) $\mathbb{h}_\nabla^{\text{pred}}$ is a widening operator.*

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

In example of Fig. 1(b), $\text{pred}(\text{DomOp}) = \{x \leq y, 4 \leq x, y \leq 100 \cdot x\}$. The refined domain operations keep the predicate $x \leq y$, which is stable among loop iterations, and hence is a loop invariant.

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”) via executable annotations. 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: For instance, if in Fig. 1(c) 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 $\mathbb{h}_{\nabla^{\text{Poly}}}^{\text{pred}}$ cannot figure out that $x \leq y$, and hence the analyzer cannot prove that the assertion is valid. Semantic hints (Sect. 6.3) will fix it.

5.1 Thresholds hints

Widening with threshold has been introduced in [2] 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 $p \leq k$, where p is some polynomial. For instance $x \in [-2, 4]$

is equivalent to $\{-x \leq 2, x \leq 4\}$. 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:

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

Lemma 2. \mathbb{h}_∇^T is: (i) a hint; and (ii) a widening.

Example 1 (Widening with thresholds). Let us consider the code snippets in Fig. 4 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 [4] for a definition of narrowing of `Intv`). In the second case, the narrowing is of no help as $\Delta([0, +\infty], \gamma([-\infty, 1000], [1002, +\infty])) = [0, +\infty]$. A widening with Thresholds $T = \{1000\}$ helps discovering the tightest loop invariant for both examples in one step as $\mathbb{h}_\nabla^T([0, 0], [1, 1]) = [0, 1000]$. \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.

The set T of thresholds is a parameter of the analyzer, which can either be provided by the user, preset to some common values (*e.g.*, $T = \{-1, 0, 1\}$), or extracted from the program text. In `Clousot`, we use a function $\text{const} \in [\text{Stm} \rightarrow \mathcal{P}(\text{int})]$ which extracts the constants appearing in the guards. We found the hint $\mathbb{h}_\nabla^{\text{const}}$ very satisfactory: (i) it helps inferring precise *numerical* loop invariants without requiring the extra iteration steps required for applying the narrowing; and (ii) it improves the precision of the analysis of code involving disequalities, *e.g.*, Fig. 4(b). A drawback 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.

6 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).

6.1 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 [7, 15]. 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 point-wise 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 $\diamond^\times((\bar{\mathbf{e}}_0^0, \bar{\mathbf{e}}_1^0) \dots (\bar{\mathbf{e}}_0^{n-1}, \bar{\mathbf{e}}_1^{n-1})) = (\bar{\diamond}_0(\bar{\mathbf{e}}_0^0 \dots \bar{\mathbf{e}}_0^{n-1}), \bar{\diamond}_1(\bar{\mathbf{e}}_1^0 \dots \bar{\mathbf{e}}_1^{n-1}))$. 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 2 (Cartesian join). Let us consider the abstract domain $Z = \text{Intv} \times \text{LT}$, where $\text{LT} = [\mathbf{Var} \rightarrow \mathcal{P}(\mathbf{Var})]$ is an abstract domain capturing the “less than” relation between *variables*. For instance, $\mathbf{x} < \mathbf{y} \wedge \mathbf{x} < \mathbf{z}$ is represented in LT by $[\mathbf{x} \mapsto \{\mathbf{y}, \mathbf{z}\}]$. The domain operations are defined as one may expect [18]. Let $\bar{\mathbf{z}}_0 = ([\mathbf{x} \mapsto [-\infty, 0]], \mathbf{y} \mapsto [1, +\infty]), [\cdot]$ and $\bar{\mathbf{z}}_1 = ([\cdot], [\mathbf{x} \mapsto \{\mathbf{y}\}])$ be two elements of Z ($[\cdot]$ denotes the empty map). Then the Cartesian join loses all the information: $\Upsilon^\times(\bar{\mathbf{z}}_0, \bar{\mathbf{z}}_1) = ([\cdot], [\cdot])$. \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 $\rho \in [\mathbf{A}^\times \rightarrow \mathbf{A}^\times]$ be a saturation (*a.k.a.* closure) procedure. Then the next lemma provides a systematic way to refine an operator \diamond^\times .

Lemma 3. *The operator $\text{h}_{\diamond^\times}^\rho$ below is a hint.*

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

Example 3 (Cartesian join, continued). The saturation of $\bar{\mathbf{z}}_0$ materializes the constraint $\mathbf{x} < \mathbf{y} : \bar{\mathbf{r}}_0 = ([\mathbf{x} \mapsto [-\infty, 0]], \mathbf{y} \mapsto [1, +\infty], [\mathbf{x} \mapsto \{\mathbf{y}\}])$, and it leaves $\bar{\mathbf{z}}_1$ unchanged. The constraint $\mathbf{x} < \mathbf{y}$ is now present in both the operands, and it is retained by the pairwise join. \square

It is worth noting that in general $\text{h}_{\diamond^\times}^\rho$ 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 [23]).

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 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.

6.2 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. 2 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:

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

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 4. $\mathbb{h}_{(\nabla, \{0,1\})}^d$ and $\mathbb{h}_{(\nabla, \{0\})}^d$ are hints and $\mathbb{h}_{(\nabla, \{0\})}^d$ is a widening.

6.3 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 5 (Computed hints). Let $\bar{e}_0, \bar{e}_1 \in \mathbf{A}$, $\Xi \in [\mathbf{A} \times \mathbf{A} \rightarrow \mathbf{A}]$ a function which returns a set of likely bounds of $\bar{e}_0 \nabla \bar{e}_1$. Then $\mathbb{h}_{\nabla}^{\Xi}$ below is a hint.

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

Computed hints are useful when the abstract join ∇ is not optimal. Otherwise, $\mathbb{h}_{\nabla}^{\Xi}$ 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 approximating the concrete union. As a consequence, no further information can be extracted from the operands. It is worth noting that in general $\mathbb{h}_{\nabla}^{\Xi}$ is not a widening. However, one can extend the arguments of the previous section to define an asymmetric hint $\mathbb{h}_{\nabla}^{\Xi}$.

Template hints Let $\text{A.range} \in [\text{Exp} \times \mathbf{A} \rightarrow \text{Intv}]$ be a function that returns the range for an expression in some abstract state, *e.g.*, it satisfies: $\forall \mathbf{E}. \forall \bar{e} \in \mathbf{A}. \text{A.range}(\mathbf{E}, \bar{e}) = [l, u] \implies \forall \sigma \in \gamma(\bar{e}). l \leq \mathbb{E}[\mathbf{E}](\sigma) \leq u$. If $\text{A.range}(\mathbf{E}, \bar{e}_i) = [l_i, u_i]$ for $i \in \{0, 1\}$, then $\gamma(\sqcup_{\text{Intv}}([l_0, u_0], [l_1, u_1]))$ is an upper bound for \mathbf{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 :

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

The main difference between $\mathbb{h}_{\overline{\mathcal{V}}}^{\Xi^P}$ 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 the counter-example of Sect. 1 with the set of templates $Oct \equiv \{\mathbf{x}_0 - \mathbf{x}_1 \mid \mathbf{x}_0, \mathbf{x}_1 \text{ are program variables}\}$. In general, template hints with Oct refine Poly so to make it as precise as Oct.

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

Fig. 5. Example requiring the use of 2D-convex hull hints

ables by: (i) projecting the abstract states on all the two-dimensional planes; and (ii) computing the planar convex hull 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. Let us consider the code in Fig. 5 from [8]. 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 [17].

6.4 Reductive hints

Intuitively, one way to improve the precision of a unary operator is to iterate its application [13]. However, an unconditional iteration may be source of unsoundness. For instance, let $- \in [\text{Intv} \rightarrow \text{Intv}]$ be the operator which applies the unary minus to an interval. In general, $\forall n \in \mathbb{N}. \bar{e} = -^{2n}(\bar{e}) \neq -^{2n+1}(\bar{e})$. 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 6 (Reductive hints). *Let $\diamond \in [\mathbb{C} \rightarrow \mathbb{C}]$ be a unary operator and $\bar{\diamond} \in [\mathbb{A} \rightarrow \mathbb{A}]$ 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.*

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 6. Abstract compilation can be used to express ψ in terms of domain operations, their compositions and state update. Lemma 6 justifies the use of local fixpoint iterations to refine the result of the

```

public BitArray(byte[] bytes) {
    Contract.Requires(bytes != null);

    this.m_array = new int[(bytes.Length + 3) / 4];
    this.m_length = bytes.Length * 8;
    int index = 0, j = 0;
    for (; (bytes.Length - j) >= 4; j+=4)
        this.m_array[index++] = (((bytes[j] & 0xff) | ((bytes[j + 1] & 0xff) << 8))
            | ((bytes[j + 2] & 0xff) << 0x10)) | ((bytes[j + 3] & 0xff) << 0x18);

    switch ((bytes.Length - j)) {
        case 1 : goto Label_00DB;
        case 2 : break;
        case 3 : this.m_array[index] = (bytes[j + 2] & 0xff) << 0x10; break;
        default: goto Label_00FC;
    }
    this.m_array[index] |= (bytes[j + 1] & 0xff) << 8;
Label_00DB:
    this.m_array[index] |= bytes[j] & 0xff;
Label_00FC:
    this.version = 0;
}

```

Fig. 6. Example of code from `mscorlib.dll`. Out of the 23 total array bound checks, `Clousot` with $\langle \text{Pnt}, \mathbb{H}_{\nabla, \nabla}^d \rangle$ validates 13, `Clousot` with $\langle \text{SubPoly}, \emptyset \rangle$ validates 6 more, and `Clousot` with $\langle \text{SubPoly}, \mathbb{H}_{\nabla}^d \rangle$ validates the remaining 4.

analysis. For instance, in the abstract domain $[\text{Var} \rightarrow \{\text{true}, \text{false}, \top, \perp\}]$ the abstract compilation of the predicate $\text{b1} == \text{b2} \wedge \text{b2} == \text{b3}$ is $\psi \equiv \lambda \mathbf{b}. (\mathbf{b}[\text{b1}, \text{b2} \mapsto \mathbf{b}(\text{b1}) \wedge \mathbf{b}(\text{b2})]) \hat{\wedge} (\mathbf{b}[\text{b2}, \text{b3} \mapsto \mathbf{b}(\text{b2}) \wedge \mathbf{b}(\text{b3})])$, where $\hat{\wedge}$ denotes the pointwise extension of \wedge . In an initial abstract state $\mathbf{b}_0 = [\text{b1}, \text{b2} \mapsto \top; \text{b3} \mapsto \text{true}]$, $\psi(\mathbf{b}_0) = [\text{b1} \mapsto \top; \text{b2}, \text{b3} \mapsto \text{true}]$ is refined by $\psi^2(\mathbf{b}_0) = [\text{b1}, \text{b2}, \text{b3} \mapsto \text{true}] = \psi^n(\mathbf{b}_0)$, $n \geq 2$.

7 Experience

We implemented hints in `Clousot`, our abstract interpretation-based static analyzer for `.Net`. `Clousot` has been designed and it is used as the static checker for the `CodeContracts` project [21]. `CodeContracts` provide a language-agnostic approach to the definition of object invariants, method preconditions and post-conditions. Contracts are specified by static methods of the `Contracts` class, *e.g.*, `Contracts.Requires(x! = null)`; specifies that the parameter `x` should be not null. More details on the specification language can be found in the documentation on the `CodeContracts` website[21]. The `Contracts` class will be shipped in the version 4.0 of the `.Net` framework [22] (at the moment of writing, in the public beta 1 phase). `Clousot` is shipped on the `DevLabs` [21] website, and it is available for free downloading for Academic use at <http://research.microsoft.com/en-us/projects/contracts/>.

Assembly Methods	P.O. Checked	Pnt		Pnt + $\mathbb{h}_{\gamma, \nabla}^d$ Slow-			Pnt + $\mathbb{h}_{\{\perp, \nabla\}^\times}^p$ Slow-		
		Valid	Time	Valid	Time	down	Valid	Time	down
<code>mscorlib</code>	17 286	14 059	3:03(0)	14 293	3:10(0)	1.0x	14 220	10:33(4)	3.3x
<code>System</code>	15 497	12 037	2:28(0)	10 321	2:36(0)	1.0x	10 143	9:43(2)	3.7x
<code>System.Web</code>	23 655	14 304	2:49(0)	13 034	2:55(0)	1.0x	13 048	8:30(0)	2.9x
<code>System.Design</code>	12 922	10 577	2:18(0)	10 135	2:21(0)	1.0x	9 947	7:39(5)	3.2x

Fig. 7. The experimental results of refining `Pnt` with die-hard hints and saturation hints. `Pnt` with die-hard hints validates 1 231 more proof obligations. `Pnt` with saturation hints are 3x slower, hit 11 timeouts (2 min), and validate 425 less accesses than \mathbb{h}^d .

`Clousot` analyzes each method `m` in isolation. It assumes the precondition of `m`, it propagates it through the body, it computes loop invariants, and it uses the inferred invariants to validate: (i) the method postcondition; (ii) the preconditions of the methods invoked by `m`; (iii) the user provided assertions; and (iv) the absence of runtime errors (*e.g.*, null pointers, array out-of-bounds, divisions by zero, negation of `MinInt` ...) and of buffer overruns [10]. When a method has no annotations, `Clousot` simply assumes the worst case scenario (*e.g.*, the parameters can assume any value compatible with their type). Orthogonally, `Clousot` can infer pre-conditions and post-conditions to help reduce the annotation burden. `Clousot` analyzes `m` incrementally. The user specifies a sequence of pairs of domains and set of hints $\langle A_0, H_0 \rangle \dots \langle A_n, H_n \rangle$. `Clousot` instantiates the abstract semantics of `m` with the abstract domain A_i refined with the hints in H_i . If it cannot discharge all the proof obligations, `Clousot` tries to discharge the remaining proof obligations using the abstract domain A_{i+1} refined with the hints H_{i+1} . We designed new numerical abstract domains, ranging from imprecise yet very fast (`Pnt`, [18]) to very precise but more expensive (`SubPoly`, [17]). In the incremental setting of `Clousot`, hints allow a very fine tuning of the precision/cost ratio. For instance, the same abstract domain can be refined with several hints: the more the hints, the more precise the analysis, but also the more expensive it is.

We report the experimental results of refining the abstract operations of the two extremes of the precision spectrum of `Clousot`'s numerical abstract domains: `Pnt` and `SubPoly`. `Pnt` is a weakly relational domain which captures properties in the form of $x \in [a, b] \wedge x < y$. `SubPoly` is a strongly relational domain which is as expressive as `Poly`, but drops some of the inference power to achieve scalability: Hints are cardinal to recover precision yet maintaining performance. We run the experiments on a Core2 Duo E6850@3.00 Ghz PC, with 4 GB of RAM, running Windows 7. We analyzed four of the main libraries of the current release of the `.Net` framework (v.3.5), available in every Windows distribution. The `mscorlib.dll` library is the core of the `.Net` framework: it contains the definitions for the `Object`, `Int32` ... types, but also common data structures such as `List`, `Dictionary`, and many other useful classes (for reflection, security ...). The `System.dll` library is a higher layer on `mscorlib.dll`. `System.Web.dll` and

`System.Design.dll` contain classes that simplify the access to the Web and the creation of user interfaces. In order to provide an uniform and repeatable test bench: (i) we considered shipped assemblies (hence with no annotations: The annotation processing is undergoing internally at Microsoft); (ii) we turned off the inference capabilities of `Clousot`; and (iii) we used `Clousot` *only* to check array creations and accesses (lower and upper bounds): the shipped assemblies do not contain annotations, so there are no contracts to check. The framework libraries contains tenths of thousands of array accesses, some of them are quite easy (*e.g.*, the sequential access of an array in a `for` loop) but others require inferring more complex relations between the array lengths and the indexes. For instance, Fig. 6 shows the constructor of the `BitArray` type (we picked it randomly from `Clousot`'s log). The `Pnt` and `SubPoly` abstract domains alone can be used to prove most of the array accesses correct, however, all the proof obligations can be discharged only using die-hard hints. One may object that the same result can be obtained using existing domains such as `Oct` or `Poly`. However, `Oct` is unable to capture the constraint `4 * m.array.Length - bytes.Length == 3`, which is necessary to prove that `index < m.array.Length`, and `Poly` suffers of a huge scalability problem, which shows up even in small code snippets like the one in Fig. 6.

Figure 7 compares die-hard hints and saturation hints when used to refine the join and widening of `Pnt`. The figure reports the analyzed assemblies, the total number of analyzed methods, the total number of proof obligations checked (*i.e.*, array creations, lower bounds, and upper bounds), the number of proof obligations validated and the analysis time for the pair-wise gathering operations and two refinements of the `Pnt` operations. The values in brackets denote the number of methods for which the analysis timed out. Time out was set to 2 minutes. Die-hard hints allow `Clousot` to validate 1 231 accesses more than the pair-wise joins at no extra cost. On the other hand, saturation hints induce an average 3x slow-down of the analysis, which causes the analysis to time out for 11 methods, and hence to validate 425 less accesses. We manually inspected the analysis logs. We found that $\langle \text{Pnt}, \mathbb{h}_{\gamma, \nabla}^d \rangle$ missed only few validations w.r.t. $\langle \text{Pnt}, \mathbb{h}_{\{\gamma, \nabla\}^*}^p \rangle$. As a consequence, the use of a saturation procedure with `Pnt` seems to be disadvantageous: the cost is too high, and the precision can be recovered by more precise abstract domains anyway. Furthermore, we checked some of the proof obligations reported as unproven or unreachable from `Clousot`. Most of the unproven conditions are caused by the lack of contracts (mainly postconditions and object invariants). However, some of the unproven conditions turned out to be real bugs, and the unreachable ones, after fixing some bug of the analyzer, were effectively dead-code.

Figure 8 focuses on the analysis of `mcorlib` using `SubPoly` refined with hints. `SubPoly` is a very expressive abstract domain (as expressive as `Poly`), whose inference precision can be fine tuned thanks to hints. 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 inter-

SubPoly		SubPoly*		Slow	SubPoly* + $\mathbb{h}_{\bar{\gamma}}^{\Xi^{Oct}}$		Slow	SubPoly* + $\mathbb{h}_{\bar{\gamma}}^{\Xi^{2DCH}}$		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. 8. The experimental results analyzing `mcorlib` with `SubPoly` and different semantic hints. `SubPoly*` denotes `SubPoly` refined with $\mathbb{h}_{\bar{\gamma}}^{\text{pred}}$ and $\mathbb{h}_{\bar{\gamma}, \nabla}^d$. Computed hints significantly slow-down the analysis, but they are needed to reach a very low false alarm ratio.

vals. 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*`, $\langle \text{SubPoly}^*, \mathbb{h}_{\bar{\gamma}}^{\Xi^{Oct}} \rangle$ and $\langle \text{SubPoly}^*, \mathbb{h}_{\bar{\gamma}}^{\Xi^{2DCH}} \rangle$ 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. 1(b) and Fig. 5 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 hints, a technique to systematically refine abstract domain *operations*. Hints allow us improving the precision of abstract operation whenever those are not complete, *e.g.*, when the underlying abstract domain is not a complete lattice (the common case in practice). We formalized hints in a relaxed abstract interpretation setting, we proved their soundness, and we distinguished between syntactic and semantics hints. We showed how some existing techniques to improve the precision of static analyses, such as widening with thresholds and reductive iterations are just instances of hints. We applied hints to the numerical abstract domains defined in our abstract interpretation-based analyzer, showing how they enable a powerful tuning of the precision/cost ratio. However, hints are not restricted to numerical domains, and they can be easily generalized to other kind of domains (for instance, for heap analysis) Future work will consider combining hints with other forms of refinement, as domain refinement, counter example-based refinement, and inference of optimal transfer functions.

References

1. R. Bagnara, P.M. Hill, and E. Zaffanella. Widening operators for powerset domains. In *VMCAI'04*, 2004.

2. 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.
3. C. Colby and P. Lee. Trace-based program analysis. In *POPL'96*. ACM Press, January 1996.
4. 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*.
5. P. Cousot and R. Cousot. Abstract interpretation and application to logic programs. *Journal of Logic Programming*, 13(2-3):103-179, July 1992.
6. P. Cousot and R. Cousot. Abstract interpretation frameworks. *Journal of Logic and Computation*, 2(4), August 1992.
7. 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.
8. P. Cousot and N. Halbwachs. Automatic discovery of linear restraints among variables of a program. In *POPL '78*.
9. M. Das, S. Lerner, and M. Seigle. Esp: Path-sensitive program verification in polynomial time. In *PLDI'02*, pages 57-68, 2002.
10. P. Ferrara, F. Logozzo, and M. A. Fähndrich. Safer unsafe code in .Net. In *OOPSLA'08*.
11. L. Gonnord and N. Halbwachs. Combining widening and acceleration in linear relation analysis. In *SAS*, pages 144-160. Springer-Verlag, 2006.
12. D. Gopan and T. W. Reps. Lookahead widening. In *CAV'06*. Springer-Verlag, August 2006.
13. P. Granger. Improving the results of static analyses programs by local decreasing iteration. In *FSTTCS'92*. Springer-Verlag.
14. B. S. Gulavani, S. Chakraborty, A. V. Nori, and S. K. Rajamani. Automatically refining abstract interpretations. In *TACAS'08*.
15. S. Gulwani, B. McCloskey, and A. Tiwari. Lifting abstract interpreters to quantified logical domains. In *POPL'08*. ACM Press, 2008.
16. M. Handjjeva and S. Tzolovski. Refining static analyses by trace-based partitioning using control flow. In *SAS '98*. Springer-Verlag, 1998.
17. V. Laviro and F. Logozzo. Subpolyhedra: a (more) scalable approach to infer linear inequalities. In *VMCAI'09*.
18. F. Logozzo and M. A. Fähndrich. Pentagons: A weakly relational abstract domain for the efficient validation of array accesses. In *SAC'08*.
19. R. Manevich, J. Field, T. A. Henzinger, G. Ramalingam, and M. Sagiv. Abstract counterexample-based refinement for powerset domains. In *Program Analysis and Compilation*, pages 273-292, 2006.
20. L. Mauborgne and X. Rival. Trace partitioning in abstract interpretation based static analyzers. In *ESOP'05*, pages 5-20, 2005.
21. Microsoft. Codecontracts tools. <http://msdn.microsoft.com/en-us/devlabs/dd491992.aspx>.
22. Microsoft. Contracts namespace. [http://msdn.microsoft.com/en-us/library/system.diagnostics.contracts\(VS.100\).aspx](http://msdn.microsoft.com/en-us/library/system.diagnostics.contracts(VS.100).aspx).
23. A. Miné. The octagon abstract domain. In *WCRE 2001*.
24. D. Monniaux. Automatic modular abstractions for linear constraints. In *POPL'09*, 2009.
25. M. Péron and N. Halbwachs. An abstract domain extending difference-bound matrices with disequality constraints. In *VMCAI'07*, 2007.
26. G. Yorsh, T. W. Reps, and S. Sagiv. Symbolically computing most-precise abstract operations for shape analysis. In *TACAS'04*. Springer-Verlag, April 2004.