



HAL
open science

Verified Functional Programming of an Abstract Interpreter

Lucas Franceschino, Jean-Pierre Talpin, David Pichardie

► **To cite this version:**

Lucas Franceschino, Jean-Pierre Talpin, David Pichardie. Verified Functional Programming of an Abstract Interpreter. SAS 2021 - 28th Static Analysis Symposium, Oct 2021, Chicago, United States. pp.1-20. hal-03342997

HAL Id: hal-03342997

<https://inria.hal.science/hal-03342997v1>

Submitted on 13 Sep 2021

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Verified Functional Programming of an Abstract Interpreter

Lucas Franceschino¹ [0000-0002-5683-0199], David Pichardie²
[0000-0002-2504-1760], and Jean-Pierre Talpin¹ [0000-0002-0556-4265]

¹ INRIA Rennes, France ² ENS Rennes, France

Abstract. Abstract interpreters are complex pieces of software: even if the abstract interpretation theory and companion algorithms are well understood, their implementations are subject to bugs, that might question the soundness of their computations.

While some formally verified abstract interpreters have been written in the past, writing and understanding them requires expertise in the use of proof assistants, and requires a non-trivial amount of interactive proofs. This paper presents a formally verified abstract interpreter fully programmed and proved correct in the F* verified programming environment. Thanks to F* refinement types and SMT prover capabilities we demonstrate a substantial saving in proof effort compared to previous works based on interactive proof assistants. Almost all the code of our implementation, proofs included, written in a functional style, are presented directly in the paper.



1 Introduction

Abstract interpretation is a theory of sound approximation. However, most of available abstract interpreters do not formally establish a relation between their algorithmic theory and implementations. Several abstract interpreters have been proven correct. The most notable one is Verasco [11], a static analyser of C programs that has been entirely written, specified and proved in the proof assistant Coq. However, understanding the implementation and proof of Verasco requires an expertise with Coq and proof assistants.

Proofs in Coq are achieved thanks to extensive use of proof scripts, that are very difficult for non expert to read. By contrast with a handwritten proof, a Coq proof can be very verbose, and often does not convey a good intuition for the idea behind a proof. Thus, writing and proving sound a static analyzer is a complex and time-consuming task: for example, Verasco requires about 17k

lines [11] of manual Coq proofs. Such an effort, however, yields the strongest guarantees and provides complete trust in the static analyser.

This paper showcases the implementation of a sound static analyser using the general-purpose functional programming language F*. Equipped with dependent types and built-in SMT solver facilities, F* provides both an OCaml-like experience and proof assistant capacities. It recently shined with the Project Everest [1], which delivered a series of formally verified, high-performance, cryptographic libraries: HACLS* [16], ValeCrypt [4] and EverCrypt [15]; that are for instance used and deployed in Mozilla Firefox. While F* can always resort to proof scripts similar to Coq's ones, most proof obligations in F* are automatically discharged by the SMT solver Z3 [9].

We present an abstract interpreter equipped with the numerical abstract domain of intervals, forward and backward analyses of expressions, widening, and syntax-directed loop iteration. This paper makes the following contributions.

- It demonstrates the ease of use of F* for verified static analysis: we implement a verified abstract interpreter, and show about 95% of its 527 lines of code (proof included) directly in the paper.
- As far as we know, it is the first time SMT techniques are used for verifying an abstract interpreter.
- We gain an order of magnitude in the number of proof lines in comparison with similar works implemented in Coq.

Related work Efforts in verified abstract interpretation are numerous [8,5,3,14], and go up to Verasco [11], a modular, real-world abstract interpreter verified in Coq. Blazy et al. [3] and Verasco follow closely the modular design of Astrée [6]; we exhibit a similar modularity on a smaller scale. However, such analysers require a non-trivial amount of mechanized proofs: in contrast, this paper shows that implementing a formally verified abstract interpreter with very little manual proofs is possible. So far, verified abstract interpreters have been focused on concretization-based formalizations. The work of Darais et al. [7] is the only one to really consider the use of Galois connections. They provide a minimalist abstract interpreter for imperative language but this interpreter seems very limited compared to ours. They use the Agda proof assistant which is comparable to Coq in terms of proof verbosity.

Overview Section 2 defines IMP, the language our abstract interpreter deals with, to which is given an operational semantics in Section 3. Then Section 4 formalizes lattices and abstract domains, while Section 5 instantiates them with the abstract domain of intervals. Section 6 derives more specific abstract domains, for numeric expressions and for memories. The latter is instantiated by Section 7, that implements an abstract weakly-relational memory. Finally, Section 8 presents the abstract interpretation of IMP statements.

The F* development is available on GitHub¹ or as supplementary material [2]. The resulting analyser is available online as a web application at <https://w95psp.github.io/verified-abstract-interpreter>.

2 IMP: a Small Imperative Language

To present our abstract interpreter, we first show the language on which it operates: IMP. It is a simple imperative language, equipped with memories represented as functions from variable names `varname` to signed integers, `intm`. This presentation lets the reader unfamiliar with F* get used to its syntax: IMP's F* definition looks like OCaml; the main difference is the explicit type signatures for constructors in algebraic data types. IMP has numeric expressions, encoded by the type `expr`, and statements `stmt`. Booleans are represented numerically: 0 represents `false`, and any other value stands for true. The enumeration `binop` equips IMP with various binary operations. The constructor `Unknown` encodes an arbitrary number. Statements in IMP are the assignment, the non-deterministic choice, the sequence and the loop.

```

type varname = | VA | VB | VC | VD   type mem τ = varname → τ
type binop   = | Plus | Minus | Mult | Eq | Lt | And | Or
type expr    = | Const: intm → expr   | Var: varname → expr
              | BinOp: binop → expr → expr → expr | Unknown
type stmt    = | Assign: varname → expr → stmt | Assume: expr → stmt
              | Seq:      stmt → stmt → stmt | Loop:  stmt → stmt
              | Choice:  stmt → stmt → stmt

```

The type `intm` is a *refinement* of the built-in F* type `ℤ`: while every integer lives in the type `ℤ`, only those that respect certain bounds live in `intm`. Numerical operations (+, - and ×) on machine integers wrap on overflow, i.e. adding one to the maximal machine integer results in the minimum machine integer. We do not give the detail of their implementation.

3 Operational Semantics

This section defines an operational semantics for IMP. It is also a good way of introducing more F* features.

We choose to formulate our semantics in terms of sets. Sets are encoded as maps from values to propositions `prop`. Those are logical statements and shouldn't be confused with booleans. Below, \subseteq quantifies over every *inhabitant* of a type: stating whether such a statement is true or false is clearly not computable. Arbitrarily complex properties can be expressed as propositions of type `prop`.

In the listing below, notice the greek letters: we use them throughout the paper. They denote implicit type arguments: for instance, below, \in works for any set `set` τ , with any type τ . F* provides the propositional operators \wedge , \vee

¹ <https://github.com/W95Psp/verified-abstract-interpreter>

and ==, in addition to boolean ones (&&, || and =). We use them below to define the union, intersection and differences of sets.

```

type set  $\tau$  =  $\tau \rightarrow \text{prop}$ 
let ( $\cap$ )  $s_0 s_1 = \lambda x \rightarrow x \in s_0 \wedge x \in s_1$ 
let ( $\cup$ ) ( $s_0 s_1$ : set  $\tau$ ): set  $\tau = \lambda x \rightarrow x \in s_0 \vee x \in s_1$ 
let ( $\subseteq$ ) ( $s_0 s_1$ : set  $\tau$ ): prop =  $\forall (x: \tau). x \in s_0 \implies x \in s_1$ 
let set_inverse (s: set intm): set intm =  $\lambda(i: \text{int}_m) \rightarrow s (-i)$ 
let ( $\in$ ) (x:  $\tau$ ) (s: set  $\tau$ ) = s x
let ( $\setminus$ )  $s_0 s_1 = \lambda v \rightarrow s_0 v \wedge \neg(s_1 v)$ 

```

To be able to work conveniently with binary operations on integers in our semantics, we define `lift_binop`, that lifts them as set operations. For example, the set `lift_binop (+) a b` (a and b being two sets of integers) corresponds to $\{va + vb \mid va \in a \wedge vb \in b\}$.

```

let lift_binop (op:  $\tau \rightarrow \tau \rightarrow \tau$ ) (a b: set  $\tau$ ): set  $\tau$ 
=  $\lambda r \rightarrow \exists (va: \tau). \exists (vb: \tau). va \in a \wedge vb \in b \wedge r == \text{op } va \ vb$ 
unfold let lift op = lift_binop (concrete_binop op)

```

The binary operations we consider are enumerated by `binop`. The function `concrete_binop` associates these syntactic operations to integer operations. For convenience, `lift` maps a `binop` to a set operation, using `lift_binop`. This function is inlined by F* directly when used because of the keyword `unfold`; intuitively `lift` behaves as a macro.

```

unfold let concrete_binop (op: binop): intm  $\rightarrow$  intm  $\rightarrow$  intm
= match op with | Plus  $\rightarrow$  nadd | Lt  $\rightarrow$  ltm | ... | Or  $\rightarrow$  orim

```

The operational semantics for expressions is given as a map from memories and expressions to sets of integers. Notice the use of both the syntax `val` and `let` for the function `osemexpr`. The `val` syntax gives `osemexpr` the type `mem \rightarrow expr \rightarrow set intm`, while the `let` declaration gives its definition. The semantics itself is uncomplicated: `Unknown` returns the set of every `intm`, a constant or a `Var` returns a singleton set. For binary operations, we lift them as set operations, and make use of recursion.

```

val osemexpr: mem  $\rightarrow$  expr  $\rightarrow$  set intm
let rec osemexpr m e =  $\lambda(i: \text{int}_m)$ 
 $\rightarrow$  match e with | Const x  $\rightarrow$  i==x | Var v  $\rightarrow$  i==m v | Unknown  $\rightarrow$  T
| BinOp op x y  $\rightarrow$  lift op (osemexpr m x) (osemexpr m y) i

```

The operational semantics for statements maps a statement and an initial memory to a set of admissible final memories. Given a statement `s`, an initial memory `mi` and a final one `mf`, `osemstmt s mi mf` (defined below) is a proposition stating whether the transition is possible.

```

val osemstmt (s: stmt): mem  $\rightarrow$  set mem
let rec osemstmt (s: stmt) (mi mf: mem)
= match s with
| Assign v e  $\rightarrow$   $\forall w. \text{if } v = w \text{ then } m_f \ v \in \text{osem}_{\text{expr}} \ m_i \ e$ 

```

```

                                else m_f w == m_i w
| Seq a b → ∃ (m_1 : mem). m_1 ∈ osem_stmt a m_i ∧ m_f ∈ osem_stmt b m_1
| Choice a b → m_f ∈ (osem_stmt a m_i ∪ osem_stmt b m_i)
| Assume e → m_i == m_f ∧ (∃ (x : int_m). x ≠ 0 ∧ x ∈ osem_expr m_i e)
| Loop a → closure (osem_stmt a) m_i m_f

```

The simplest operation is the assignment of a variable v to an expression e : the transition is allowed if every variable but v in m_i and m_f is equal and the final value of v matches with the semantics of e . Assuming that an expression is true amounts to require the initial memory to be such that at least a non-zero integer (that is, the encoding of `true`) belongs to $\text{osem_expr } m_i \ e$. The statement `Seq a b` starting from the initial memory m_i admits m_f as a final memory when there exists (i) a transition from m_i to an intermediate memory m_1 with statement a and (ii) a transition from m_1 to m_f with statement b . The operational semantics for a loop is defined as the reflexive transitive closure of the semantics of its body. The `closure` function computes such a closure, and is provided by F^* 's standard library.

4 Abstract Domains

Our abstract interpreter is parametrized over relational domains. We instantiate it later with a weakly-relational [6] memory. This section defines lattices and abstract domains. Such structures are a natural fit for typeclasses [13], which allow for ad hoc polymorphism. In our case, it means that we can have one abstraction for lattices for instance, and then instantiate this abstraction with implementations for, say, sets of integers, then intervals, etc. Typeclasses can be seen as record types with dedicated dependency inference. Below, we define the typeclass `lattice`: defining an instance for a given type equips this type with a lattice structure.

Refinement types Below, the syntax $x:\tau\{p \ x\}$ denotes the type whose inhabitants both belong to τ and satisfy the predicate p . For example, the inhabitant of the type `bot : ℕ{∀(n:ℕ). bot ≤ n}` is 0: it is the (only) smallest natural number. To typecheck $x:\tau$, F^* collects the *proof obligations* implied by "x has the type τ ", and tries to discharge them with the help of the SMT solver. If the SMT solver is able to deal with the proof obligations, then $x:\tau$ typechecks. In the case of "0 is of type `bot : ℕ{∀(n:ℕ). bot ≤ n}`", the proof obligation is $\forall(n:\mathbb{N}). 0 \leq n$.

Below, most of the types of the fields from the record type `lattice` are refined. Typechecking i against the type `lattice` τ yields a proof obligation asking (among other things) for `i.join` to go up in the lattice and for `bottom` to be a lower bound. Thus, if " i has type `lattice` τ " typechecks, it means there exists a proof that the properties written as refinements in `lattice`'s definition hold on i . We found convenient to let `bottom` represent unreachable states. Note `lattice` is under-specified, i.e. it doesn't require `join` to be provably a least upper bound, since such a property plays no role in our proof of soundness. This choice follows Blazy and al. [3].

```

class lattice  $\tau$  = { corder: order  $\tau$ 
  ; join:  $x:\tau \rightarrow y:\tau \rightarrow r:\tau$  {corder  $x$   $r$   $\wedge$  corder  $y$   $r$ }
  ; meet:  $x:\tau \rightarrow y:\tau \rightarrow r:\tau$  {corder  $r$   $x$   $\wedge$  corder  $r$   $y$ }
  ; bottom: bot: $\tau$ { $\forall x$ . corder bot  $x$ }; top: top: $\tau$ { $\forall x$ . corder  $x$  top}}

```

For our purpose, we need to define what an abstract domain is. In our setting, we consider concrete domains with powerset structure. The typeclass `adom` encodes them: it is parametrized by a type τ of abstract values. For instance, consider `itv` the type for intervals: `adom itv` would be the type inhabited by correct abstract domains for intervals.

Implementing an abstract domain amounts to implementing the following fields: (i) `c`, that represents the type to which abstract values τ concretizes; (ii) `adomlat`, a lattice for τ ; (iii) `widen`, a widening operator; (iv) γ , a monotonic concretization function from τ to `set c`; (v) `order_measure`, a measure ensuring the abstract domain doesn't admit infinite increasing chains, so that termination is provable for fixpoint iterations; (vi) `meetlaw`, that requires `meet` to be a correct approximation of set intersection; (vii) `toplaw` and `botlaw`, that ensure the lattice's bottom concretization matches with the empty set, and similarly for `top`.

```

class adom  $\tau$  = { c: Type; adomlat: lattice  $\tau$ 
  ;  $\gamma$ : ( $\gamma$ : ( $\tau \rightarrow$  set  $c$ ) { $\forall (x\ y:\tau)$ . corder  $x$   $y \implies (\gamma\ x \subseteq \gamma\ y)$ })
  ; widen:  $x:\tau \rightarrow y:\tau \rightarrow r:\tau$  {corder  $x$   $r$   $\wedge$  corder  $y$   $r$ }
  ; order_measure: measure adomlat.corder
  ; meetlaw:  $x:\tau \rightarrow y:\tau \rightarrow$  Lemma (( $\gamma\ x \cap \gamma\ y \subseteq \gamma$  (meet  $x\ y$ ))
  ; botlaw: unit  $\rightarrow$  Lemma ( $\forall (x:c)$ .  $\sim(x \in \gamma$  bottom))
  ; toplaw: unit  $\rightarrow$  Lemma ( $\forall (x:c)$ .  $x \in \gamma$  top)}

```

Notice the refinement types: we require for instance the monotony of γ . Every single instance for `adom` will be checked against these specifications. No instance of `adom` where γ is not monotonic can exist. With a proposition `p`, the `Lemma p` syntax signals a function whose outcome is computationally irrelevant, since it simply produces `()`, the inhabitant of the type `unit`. However, it does not produce an arbitrary `unit`: it produces an inhabitant of `_:unit {p}`, that is, the type `unit` refined with the goal `p` of the lemma itself.

For praticity, we define some infix operators for `adomlat` functions. The syntax `{|...|}` lets one formulate typeclass constraints: for example, `(\sqsubseteq)` below ask `F*` to resolve an instance of the typeclass `adom` for the type τ , and name it `l`. Below, `(\sqcap)` instantiates the lemma `meetlaw` explicitly: `meetlaw x y` is a unit value that carries a proof in the type system.

```

let ( $\sqsubseteq$ ) {|l:adom  $\tau$ |} = l.adomlat.corder
let ( $\sqcup$ ) {|l:adom  $\tau$ |} (x y: $\tau$ ): r: $\tau$  { corder  $x$   $r$   $\wedge$  corder  $y$   $r$ 
   $\wedge (\gamma\ x \cup \gamma\ y) \subseteq \gamma\ r$  } = join x y
let ( $\sqcap$ ) {|l:adom  $\tau$ |} (x y: $\tau$ ): r: $\tau$  { corder  $r$   $x$   $\wedge$  corder  $r$   $y$ 
   $\wedge (\gamma\ x \cap \gamma\ y) \subseteq \gamma\ r$  }
  = let _ = meetlaw x y in meet x y

```

Lemmas are functions that produce refined `unit` values carrying proofs. Below, given an abstract domain `i`, and two abstract values `x` and `y`, `join_lemma i x y`

is a proof concerning i , x and y . Such an instantiation can be manual (i.e. below, `i.toplaw ()` in `top_lemma`), or automatic. The automatic instantiation of a lemma is decided by the SMT solver. Below, we make use of the `SMTPat` syntax, that allows us to give the SMT solver a list of patterns. Whenever the SMT solver matches a pattern from the list, it instantiates the lemma in stake. The lemma `join_lemma` below states that the union of the concretization of two abstract values x and y is below the concretization of the abstract join of x and y . This is true because of γ 's monotony: we help a bit the SMT solver by giving a hint with `assert`. This lemma is instantiated each time a proof goal contains $x \sqsubseteq y$.

Because of a technical limitation, we cannot write SMT patterns directly in the `meetlaw`, `botlaw` and `toplaw` fields of the class `adom`: thus, below we reformulate them.

```
let top_lemma (i: adom  $\tau$ )      (let bot_lemma, meet_lemma = ...)
  : Lemma ( $\forall (x: i.c). x \in i.\gamma i.adom_{lat}.top$ )
    [SMTPat (i. $\gamma$  i.adomlat.top)] = i.toplaw ()
let join_lemma (i: adom  $\tau$ ) (x y:  $\tau$ )
  : Lemma ((i. $\gamma$  x  $\cup$  i. $\gamma$  y)  $\subseteq$  i. $\gamma$  (i.adomlat.join x y))
    [SMTPat (i.adomlat.join x y)]
  = let r = i.adomlat.join x y in assert ( $\gamma$  x  $\subseteq$   $\gamma$  r  $\wedge$   $\gamma$  y  $\subseteq$   $\gamma$  r)
```

5 An Example of Abstract Domain: Intervals

Until now, the F^* code we presented was mostly specificational. This section presents the abstract domain of intervals, and thus shows how proof obligations are dealt with in F^* . Below, the type `itv'` is a dependent tuple: the refinement type on its right-hand side component up depends on `low`. If a pair (x, y) is of type `itv'`, we have a proof that $x \leq y$.

```
type itv' = low:intm & up:intm {low $\leq$ up}    type itv = withbot itv'
```

The machine integers being finite, `itv'` naturally has a top element. However, `itv'` cannot represent the empty set of integers, whence `itv`, that adds an explicit bottom element using `withbot`. The syntax `Val?` returns true when a value is not `Bot`. For convenience, `mk` makes an interval out of two numbers, and `itvcard` computes the cardinality of an interval. We use it later to define a measure for intervals. `inbounds x` holds when $x:\mathbb{Z}$ fits machine integer bounds.

```
type withbot (a: Type) = | Val: v:a  $\rightarrow$  withbot a | Bot
let mk (x y:  $\mathbb{Z}$ ): itv = if inbounds x && inbounds y && x  $\leq$  y
  then Val (x,y) else Bot
let itvcard (i:itv): $\mathbb{N}$  = match i with | Bot  $\rightarrow$  0 | Val i  $\rightarrow$  dsnd i - dfst i + 1
```

Below, `latitv` is an instance of the typeclass `lattice` for intervals: intervals are ordered by inclusion, the `meet` and `join` operations consist in unwrapping `withbot`, then playing with bounds. `latitv` is of type `lattice itv`: it means for

instance that we have the proof that the join and meet operators respect the order `latitv.corder`, as stated in the definition of `lattice`. Note that here, not a single line of proof is required: F* transparently builds up proof obligations, and asks the SMT to discharge them, that does so automatically.

```
instance latitv: lattice itv =
  { corder = withbotord #itv' (λ(a,b) (c,d) → a>c && b<d)
  ; join = (λ(i j: itv) → match i, j with
    | Bot, k | k, Bot → k
    | Val (a,b), Val (c,d) → Val (min a c, max b d))
  ; meet = (λ(x y: itv) → match x, y with
    | Val (a,b), Val (c,d) → mk (max a c) (min b d)
    | _ → Bot); bottom = Bot; top = mk minintm maxintm }
```

Such automation is possible even with more complicated definitions: for instance, below we define the classical widening with thresholds. Without a single line of proof, `widen` is shown as respecting the order `corder`.

```
let thresholds: list intm = [minintm; -64; -32; -16; -8; -4; 4; 8; 16; 32... ]
let widen_bound_r (b: intm): (r: intm {r>b ∨ b=maxintm}) =
  if b=maxintm then b else find' (λ(u: intm) → u>b) thresholds
let widen_bound_l (b: intm): (r: intm {r<b ∨ b=minintm}) =
  if b=minintm then b else find' (λ(u: intm) → u<b) (rev thresholds)
let widen (i j: itv): r: itv {corder i r ∧ corder j r}
= match i, j with | Bot, x | x, Bot → x
  | Val (a,b), Val (c,d) → Val ( (if a≤c then a else widen_bound_l c)
    , (if b≥d then b else widen_bound_r d) )
```

Similarly, turning `itv` into an abstract domain requires no proof effort. Below `itvadom` explains that intervals concretize to machine integers (`c = intm`), how it does so (with $\gamma = \text{itv}_\gamma$), and which lattice is associated with the abstract domain (`adomlat = latitv`). As explained previously, the proof of a proposition p in F* can be encoded as an inhabitant of a refinement of `unit`, whence the "empty" lambdas: we let the SMT solver figure out the proof on its own.

```
let itvγ: itv → set intm = withbotγ (λ(i: itv') x → dfst i ≤ x ∧ x ≤ dsnd i)
instance itvadom: adom itv = { c = intm ; adomlat = latitv; γ = itvγ
  ; meetlaw = (λ_ _ → ()); botlaw = (λ_ → ()); toplaw = (λ_ → ())
  ; widen = widen ; order_measure = {f=itvcard; max=sizeintm}}
```

5.1 Forward Binary Operations on Intervals

Most of the binary operations on intervals can be written and shown correct without any proof. Our operators handle machine integer overflowing: for instance, `add_overflows` returns a boolean indicating whether the addition of two integers overflows, solely by performing machine integer operations. The refinement of `add_overflows` states that the returned boolean `r` should be true if and only if the addition in \mathbb{Z} differs from the one in `intm`. The correctness

of `itvadd` is specified as a refinement: the set of the additions between the concretized values from the input intervals is to be included in the concretization of the abstract addition. Its implementation is very simple, and its correctness proved automatically.

```

let add_overflows (a b: intm)
  : (r: bool {r  $\iff$  int_arith.nadd a b  $\neq$  int_m_arith.nadd a b})
  = ((b < 0) = (a < 0)) && abs a > maxintm - abs b
let itvadd (x y: itv): (r: itv {( $\gamma$  x +  $\gamma$  y)  $\subseteq$   $\gamma$  r})
  = match x, y with | Val (a, b), Val (c, d)
    → if add_overflows a c || add_overflows b d
      then top else Val (a + c, b + d) | _  $\rightarrow$  Bot

```

However the SMT solver sometimes misses some necessary lemmas. In such cases, we can either guide the SMT solver by discriminating cases and inserting hints, or go fully manual with a tactic system à la Coq. Below, the `assert` uses tactics: everything within the parenthesis following the `by` keyword is a computation that manipulates proof goals. Our aim is to prove that subtracting two numerical sets a and b is equivalent to adding a with the inverse of b .

Unfortunately, due to the nature of `lift_binop`, this yields existential quantifications which are difficult for the SMT solver to deal with. After normalizing our goal (with `compute ()`), and dealing with quantifiers and implications (`forall_intro`, `implies_intro` and `elim_exists`), we are left with $\exists y. b \ (-y) \wedge r=x+y$ knowing $b \ z \wedge r=x-z$ given some z as an hypothesis. Eliminating $\exists y$ with $-z$ is enough to complete the proof.

We sadly had to prove that (not too complicated) fact by hand. This however shows the power of F^* . Its type system is very expressive: one can state arbitrarily mathematically hard propositions (for which automation is hopeless). In such cases, one can always resort to Coq-like manual proving to handle hard proofs.

```

let set_inverse (s: set intm): set intm =  $\lambda$ (i: intm)  $\rightarrow$  s (-i)
let lemmainv (a b: set intm)
  : Lemma ((a-b)  $\subseteq$  (a+set_inverse b)) [SMTPat (a+set_inverse b)]
  = assert ((a-b)  $\subseteq$  (a+set_inverse b)) by ( compute ();
    let _ = forall_intro () in let p0 = implies_intro () in
    let witX,p1 = elim_exists (binder_to_term p0) in
    let witY,p1 = elim_exists (binder_to_term p1) in
    let z:  $\mathbb{Z}$  = unquote (binder_to_term witY) in
    witness witX; witness (quote (-z)))

```

Notice the SMT pattern: the lemma `lemmainv` will be instantiated each time the SMT deals with an addition involving an inverse. Defining the subtraction `itvsub` is a breeze: it simply performs an interval addition and an interval inversion. Here, no need for a single line of proof for its correctness (expressed as a refinement).

```

let itvinv (i: itv): (r: itv {set_inverse ( $\gamma$  i)  $\subseteq$   $\gamma$  r})
  = match i with | Val (lower, upper)  $\rightarrow$  Val (-upper, -lower) | _  $\rightarrow$  i
let itvsub (x y: itv): (r: itv {( $\gamma$  x -  $\gamma$  y)  $\subseteq$   $\gamma$  r}) = itvadd x (itvinv y)

```

Proving multiplication sound on intervals requires a lemma which is not inferred automatically:

$$\forall x \in [a, b], y \in [b, c]. [\min(ac, ad, bc, bd), \max(ac, ad, bc, bd)]$$

In that case, decomposing that latter lemma into sublemmas `lemma_min` and `lemma_mul` is enough. Apart from this lemma, `itv_mul` is free of any proof term.

```

let lemma_min (a b c d : ℤ) (x : ℤ{a ≤ x ∧ x ≤ b}) (y : ℤ{c ≤ y ∧ y ≤ d})
  : Lemma (x*y ≥ a*c ∨ x*y ≥ a*d ∨ x*y ≥ b*c ∨ x*y ≥ b*d) = ()
unfold let in_btw (x : ℤ) (l u : ℤ) = l ≤ u ∧ x ≥ l ∧ x ≤ u
let lemma_mul (a b c d x y : ℤ)
  : Lemma (requires in_btw x a b ∧ in_btw y c d)
    (ensures x*y ≥ (a*c) `min` (a*d) `min` (b*c) `min` (b*d)
      ∧ x*y ≤ (a*c) `max` (a*d) `max` (b*c) `max` (b*d))
  [SMTPat (x*y); SMTPat (a*c); SMTPat (b * d)]
  = lemma_min a b c d x y; lemma_min (-b) (-a) c d (-x) y

let mul_overflows (ab:int_m):(r:bool{r≠inbounds (int_arith.n_mul a b)})
  = a ≠ 0 && abs b > max_int_m `div_m` (abs a)
let itv_mul (x y : itv): r:itv {(γ x × γ y) ⊆ γ r}
  = match x, y with
  | Val (a, b), Val (c, d) →
    let l = (a*c) `min` (a*d) `min` (b*c) `min` (b*d) in
    let r = (a*c) `max` (a*d) `max` (b*c) `max` (b*d) in
    if mul_overflows a c || mul_overflows a d
    || mul_overflows b c || mul_overflows b d
    then top else Val (l, r)
  | _ → Bot

```

The forward boolean operators for intervals require no proof at all; here we only give their type signatures. A function of interest is `itv_as_bool`: it returns `TT` when an interval does not contain 0, `FF` when it is the singleton 0, `Unk` otherwise.

```

let β (x : int_m) : itv = mk x x
let itv_eq (x y : itv) : r : itv {(γ x `neq` γ y) ⊆ γ r} = ... let itv_1t = ...
let itv_cγ (i : itv) (x : int_m) : r : bool {r ⇔ itv_γ i x} = ...
let itv_as_bool (x : itv) : ubool // with type ubool = |Unk|TT|FF
  = if β 0=x || Bot?x then FF else if itv_cγ x 0 then Unk else TT
let itv_andi (x y : itv) : (r : itv {(γ x `nand` γ y) ⊆ γ r})
  = match itv_as_bool x, itv_as_bool y with
  | TT, TT → β 1 | FF, _ | _, FF → β 0 | _, _ → mk 0 1
let itv_ori (x y : itv) : (r : itv {(γ x `nor` γ y) ⊆ γ r}) = ...

```

5.2 Backward Operators

While a forward analysis for expressions is essential, another powerful analysis can be made thanks to backward operators. Typically, it aims at extracting

information from a test, and at refining the abstract values involved in this test, so that we gain in precision on those abstract values. Given a concrete binary operator \oplus , we define $\overleftarrow{\oplus}$ its abstract backward counterpart. Let three intervals $x^\#, y^\#,$ and $r^\#$. $\overleftarrow{\oplus} x^\# y^\# r^\#$ tries to find the most precise intervals $x^{\#\#}$ and $y^{\#\#}$ supposing $\gamma x^\# \oplus \gamma y^\# \subseteq \gamma r^\#$. The soundness of $\overleftarrow{\oplus} x^\# y^\# r^\#$ can be formulated as below. We later generalize this notion of soundness with the type $\text{sound}_{\delta p}$, which is indexed by an abstract domain and a binary operation.

```

let  $x^{\#\#}, y^{\#\#} = (\overleftarrow{\oplus}) x^\# y^\# r^\#$  in
   $\forall x y. (x \in \gamma x^\# \wedge y \in \gamma y^\# \wedge \text{op } x y \in \gamma r^\#)$ 
     $\implies (x \in \gamma x^{\#\#} \wedge y \in \gamma y^{\#\#})$ 

```

As the reader will discover in the rest of this section, this statement of soundness is proved entirely automatically against each and every backward operator for the interval domain. For op a concrete operator, $\text{sound}_{\delta p} \text{ itv } \text{op}$ is inhabited by sound backward operators for op in the domain of intervals. If one shows that $\overleftarrow{\oplus}$ is of type $\text{sound}_{\delta p} \text{ itv } (+)$, it means exactly that $\overleftarrow{\oplus}$ is a sound backward binary interval operator for $(+)$. The rest of the listing shows how light in proof and OCaml-looking the backward operations are. Below, we explain how $\overleftarrow{\text{lt}}$ works: it is a bit complicated because it hides a " $\overleftarrow{\text{ge}}$ " operator.

```

let  $\overleftarrow{\text{add}}$ :  $\text{sound}_{\delta p} \text{ itv } n_{\text{add}} = \lambda x y r \rightarrow x \sqcap (r-y), y \sqcap (r-x)$ 
let  $\overleftarrow{\text{sub}}$ :  $\text{sound}_{\delta p} \text{ itv } n_{\text{sub}} = \lambda x y r \rightarrow x \sqcap (r+y), y \sqcap (x-r)$ 
let  $\overleftarrow{\text{mul}}$ :  $\text{sound}_{\delta p} \text{ itv } n_{\text{mul}} = \lambda x y r \rightarrow$ 
  let  $h$  ( $i j$ :itv) = (if  $j=\beta 1$  then  $i \sqcap r$  else  $i$ ) in  $h x y, h y x$ 
let  $\overleftarrow{\text{eq}}$ :  $\text{sound}_{\delta p} \text{ itv } n_{\text{eq}}$ 
  =  $\lambda x y r \rightarrow \text{match itv\_as\_bool } r \text{ with } | \text{TT} \rightarrow x \sqcap y, x \sqcap y \mid \_ \rightarrow x, y$ 
let  $(\setminus)$  ( $x y$ : itv): ( $r$ : itv  $\{(\gamma x \setminus \gamma y) \subseteq \gamma r\}$ ) =...
let  $\overleftarrow{\text{and}}$ :  $\text{sound}_{\delta p} \text{ itv } n_{\text{and}}$ 
  =  $\lambda x y r \rightarrow \text{match itv\_as\_bool } r, \text{itv\_as\_bool } x, \text{itv\_as\_bool } y \text{ with}$ 
     $| \text{FF}, \text{TT}, \_ \rightarrow x, y \sqcap \beta 0 \quad | \text{FF}, \_, \text{TT} \rightarrow x \sqcap \beta 0, y$ 
     $| \text{TT}, \_, \_ \rightarrow x \setminus \beta 0, y \setminus \beta 0 \quad | \_ \rightarrow x, y$ 
let  $\overleftarrow{\text{or}}$ :  $\text{sound}_{\delta p} \text{ itv } n_{\text{or}}$ 
  =  $\lambda x y r \rightarrow \text{match itv\_as\_bool } r, \text{itv\_as\_bool } x, \text{itv\_as\_bool } y \text{ with}$ 
     $| \text{TT}, \text{FF}, \text{Unk} \mid \text{TT}, \text{FF}, \text{FF} \rightarrow x, y \setminus \beta 0 \mid \text{TT}, \text{Unk}, \text{FF} \rightarrow x \setminus \beta 0, y$ 
     $| \text{FF}, \_, \text{TT} \mid \text{FF}, \text{TT}, \_ \rightarrow x \sqcap \beta 0, y \sqcap \beta 0 \mid \_ \rightarrow x, y$ 

```

Let us look at $\overleftarrow{\text{lt}}$. Knowing whether $x < y$ holds, $\overleftarrow{\text{lt}}$ helps us refining x and y to more precise intervals. Let x be the interval $[0; \text{max}_{\text{int}_m}]$, y be $[5; 15]$ and r be $[0; 0]$. Since the singleton $[0; 0]$ represents **false**, $\overleftarrow{\text{lt}} x y r$ aims at refining x and y knowing that $x < y$ doesn't hold, that is, knowing $x \geq y$. In this case, $\overleftarrow{\text{lt}}$ finds $x' = [5; \text{max}_{\text{int}_m}]$ and $y' = [5; 15]$. Indeed, when r is $[0; 0]$, $\text{itv_as_bool } r$ equals to **FF**. Then we rewrite $\neg(x < y)$ either as $y < x + 1$ (when x is incrementable) or as $y - 1 < x$. In our case, x 's upper bound is $\text{max}_{\text{int}_m}$ (the biggest int_m): x is not incrementable. Thus we rewrite $\neg([0; \text{max}_{\text{int}_m}] < [5; 15])$ as $[6; 16] < [0; \text{max}_{\text{int}_m}]$.

Despite of these different case handling, the implementation of $\overleftarrow{\text{It}}$ required no proof: the SMT solver takes care of everything automatically.

```

let  $\overleftarrow{\text{It}}_{\text{true}}$  (x y: itv)
  = match x, y with | Bot, _ | _, Bot → x,y
  | Val(a,b), Val(c,d) → mk a (min b (d-1)), mk (max (a+1) c) d
let decrementable i=Val?i&&dfst(Val?.v i)>minintm let incr.=...
let  $\overleftarrow{\text{It}}$ : sound $\delta_p$  itv nIt
  =  $\lambda$ x y r → match itv_as_bool r with | TT →  $\overleftarrow{\text{It}}_{\text{true}}$  x y
  | FF → if incrementable x //  $x < y \iff y > x+1$ 
      then let ry, rx =  $\overleftarrow{\text{It}}_{\text{true}}$  y (itvadd x ( $\beta$  1)) in
          itvsub rx ( $\beta$  1), ry
      else if decrementable y //  $x < y \iff y-1 > x$ 
          then let ry, rx =  $\overleftarrow{\text{It}}_{\text{true}}$  (itvsub y ( $\beta$  1)) x in
              rx, itvadd ry ( $\beta$  1)
          else x,y | _ → x, y

```

6 Specialized Abstract Domains

Abstract domains are defined in Section 4 as lattices equipped with a sound concretization operation. Our abstract interpreter analyses IMP programs: its expressions are numerical, and IMP is equipped with a memory. Thus, this section defines two specialized abstract domains: one for numerical abstractions, and another one for memory abstractions.

6.1 Numerical Abstract Domains

In the section 5.2, we explain what a sound backward operator is in the case of the abstract domain of intervals. There, we mention a more generic type sound_{δ_p} that states soundness for such operators in the context of any abstract domain. We present its definition below:

```

type sound $\delta_p$  (a:Type) {l:adom a} (op:l.c→l.c→l.c)
  =  $\delta_p$ : (a → a → a → (a & a)) {
     $\forall$  (x# y# r#: a). let x##, y## =  $\overleftarrow{\delta_p}$  x# y# r# in
      ( $\forall$  (x y: l.c). (x  $\in$   $\gamma$  x#  $\wedge$  y  $\in$   $\gamma$  y#  $\wedge$  op x y  $\in$   $\gamma$  r#)
         $\implies$  (x  $\in$   $\gamma$  x##  $\wedge$  y  $\in$   $\gamma$  y##))}

```

We define the specialized typeclass num_{adom} for abstract domains that concretize to machine integers. A type that implements an instance of num_{adom} should also have an instance of adom , with int_m as concrete type. Whence the fields na_{adom} , and adom_{num} . Moreover, we require a computable concretization function cgamma , that is, a function that maps abstract values to computable sets of machine integers: $\text{int}_m \rightarrow \text{bool}$. The β operator lifts a concrete value in the abstract world. We also require the abstract domain to provide both sound forward and backward operator for every syntactic operator of type binop presented in Section 2.

The function `abstract_binop` maps an operator `op` of type `binop` to a sound forward abstract operator. Its soundness is encoded as a refinement. Similarly, `abstract_binop` maps a `binop` to a corresponding sound backward operator. To ease backward analysis, `gt0` and `lt0` are abstractions for non-null positive and negative integers.

```
class num_adom (a: Type) =
{ na_adom: adom a; adom_num: squash (na_adom.c == int_m)
; cgamma: x#:a → x:int_m → b:bool {b ↔ x ∈ γ x#}
; abstract_binop: op:_ → i:a → j:a → r:a {lift op (γ i) (γ j) ⊆ γ r}
; abstract_binop: (op: binop) → soundop a (concrete_binop op)
; gt0: x#:a {∀(x:int_m). x>0 ⇒ x ∈ γ x#}
; lt0: x#:a {∀(x:int_m). x<0 ⇒ x ∈ γ x#}; β: x:int_m → r:a {x ∈ γ r} }
```

For a proposition `p`, the F* standard library defines `squash p` as the type `_:unit{p}`, that is, a refinement of the unit type. This can be seen as a lemma with no parameter.

Instance for intervals The section 5 defines everything required by `num_adom`, thus below we give an instance of the typeclass `num_adom` for intervals.

```
instance itv_num_adom: num_adom itv = {
  na_adom = solve; adom_num = (); cgamma = itv_cγ; β = (λ x → β x);
  abstract_binop = (function | Plus → itv_add ... | Or → itv_ori);
  abstract_binop = (function | Plus → add ... | Or → or );
  lt0 = (mk minint_m (-1)); gt0 = (mk ( 1) maxint_m) }
```

6.2 Memory Abstract Domains

From the perspective of IMP statements, an abstract domain for abstract memories is fairly simple. An abstract memory should be equipped with two operations: assignment and assumption. Those are directly related to their syntactic counterpart `Assume` and `Assign`. Thus, `mem_adom` has a field `assume_` and a field `assign`. The correctness of these operations are elegantly encoded as refinement types.

Let us explain the refinement of `assume_`: let `m0#` an abstract memory, and `e` an expression. For every concrete memory `m0` abstracted by `m0#`, the set of acceptable final memories `osemstmt (Assume e) m0` should be abstracted by `assume_ m0# e`.

```
class mem_adom μ = { ma_adom: adom μ; ma_mem: squash (ma_adom.c == mem);
  assume_: m0#:μ → e:expr → m1#:μ
    {∀ (m0: mem {m0 ∈ γ m0#}). osemstmt (Assume e) m0 ⊆ γ m1#};
  assign: m0#:μ → v:varname → e:expr → m1#:μ
    {∀ (m0: mem {m0 ∈ γ m0#}). osemstmt (Assign v e) m0 ⊆ γ m1#}}
```

7 A Weakly-Relational Abstract Memory

In this section, we define a weakly-relational abstract memory. This abstraction is said weakly-relational because the entrance of an empty abstract value in the map systematically launches a reduction of the whole map to **Bot**. Below we define an abstract memory (**amem**) as either an unreachable state (**Bot**), or a mapping (**map** τ) from **varname** to abstract values τ . The mappings **map** τ are equipped with the utility functions **mapi**, **map₁**, **map₂** and **fold**.

```

type map  $\tau$  = ... type amem  $\tau$  = withbot (map  $\tau$ )
let get': map  $\tau$   $\rightarrow$  varname  $\rightarrow$   $\tau$  = ... let fold: ( $\tau \rightarrow \tau \rightarrow \tau$ )  $\rightarrow$  map  $\tau \rightarrow \tau$  = ...
let mapi: (varname  $\rightarrow$   $\tau \rightarrow \beta$ )  $\rightarrow$  map  $\tau \rightarrow$  map  $\beta$  = ...
let map1: ( $\tau \rightarrow \beta$ )  $\rightarrow$  map  $\tau \rightarrow$  map  $\beta$  =  $\lambda f \rightarrow$  mapi ( $\lambda\_ \rightarrow f$ )
let map2: ( $\tau \rightarrow \beta \rightarrow \gamma$ )  $\rightarrow$  map  $\tau \rightarrow$  map  $\beta \rightarrow$  map  $\gamma$  = ...

```

A lattice structure The listing below presents **amem** instances for the type-classes **order**, **lattice** and **mem_{adom}**. Once again, the various constraints imposed by these different typeclasses are discharged automatically by the SMT solver.

```

let amem_update (k: varname) (v:  $\tau$ ) (m: amem  $\tau$ ): amem  $\tau$ 
  = match m with | Bot  $\rightarrow$  Bot
    | Val m  $\rightarrow$  Val (mapi ( $\lambda k' v' \rightarrow$  if k'=k then v else v') m)
instance amemlat {l: adom  $\tau$  |}: lattice (amem  $\tau$ ) =
  { corder = withbotord ( $\lambda m_0 m_1 \rightarrow$  fold (&&) (map2 corder m0 m1))
    ; join = ( $\lambda x y \rightarrow$  match x, y with
      | Val x, Val y  $\rightarrow$  Val (map2 join x y) | m, Bot | _, m  $\rightarrow$  m)
    ; meet = ( $\lambda x y \rightarrow$  match x, y with
      | Val x, Val y  $\rightarrow$ 
        let m = map2 ( $\cap$ ) x y in
        if fold ( $\cap$ ) (mapi ( $\lambda\_ v \rightarrow$  l.adomlat.corder v bottom) m)
        then Bot else Val m
      | _  $\rightarrow$  Bot); bottom = Bot; top = ... }
instance amemadom {l: adom  $\tau$  |}: adom (amem  $\tau$ ) = { c = mem' l.c
  ; adomlat=solve; meetlat=( $\lambda\_ \rightarrow$ ()); toplat=( $\lambda\_ \rightarrow$ ()); botlat=( $\lambda\_ \rightarrow$ ())
  ;  $\gamma$  = withbot $\gamma$  ( $\lambda m^\# m \rightarrow$  fold ( $\wedge$ ) (mapi ( $\lambda v x \rightarrow$  m v  $\in \gamma$  x) m#))
  ; widen = ( $\lambda x y \rightarrow$  match x, y with
    | Val x, Val y  $\rightarrow$  Val (map2 widen x y) | m, Bot | _, m  $\rightarrow$  m)
  ; order_measure = let {max; f} = l.order_measure in
    { f = (function | Bot  $\rightarrow$  0 | Val m#  $\rightarrow$  1 + fold (+) (map1 f m#))
      ; max = 1 + max  $\times$  4 } }

```

The rest of this section defines a **mem_{adom}** instance for our memories **amem**. The typeclass **mem_{adom}** is an essential piece in our abstract interpreter: it provides the abstract operations for handling assumes and assignments.

Forward expression analysis We define **asem_{expr}**, mapping expressions to abstract values of type τ . It is defined for any abstract domain, whence the typeclass

argument $\{|\text{num}_{\text{adom}} \tau|\}$. The abstract interpretation of an expression e given $m_0^\#$ an initial memory is defined below as $\text{asem}_{\text{expr}} m_0^\# e$. It is specified via a refinement type to be a sound abstraction of e 's operational semantics $\text{osem}_{\text{expr}} m_0 e$. This function leverages the operators from the different typeclasses for which we defined instances just above. $\beta: \text{int}_m \rightarrow \tau$ and $\text{abstract_binop}: \text{binop} \rightarrow \dots$ come from num_{adom} , while $\text{top}: \tau$ comes from lattice .

```

val get: m: amem  $\tau$  {Val? m}  $\rightarrow$  varname  $\rightarrow$   $\tau$     let get (Val m) = get' m
let rec asem_expr {|\text{num}_{\text{adom}} \tau|\} (m_0^\#: amem  $\tau$ ) (e: expr)
: (r:  $\tau$  {  $\forall$  (m_0: mem). m_0  $\in$   $\gamma$  m_0^\#  $\implies$   $\text{osem}_{\text{expr}} m_0 e \subseteq \gamma r$  })
= if m_0^\#  $\sqsubseteq$  bottom then bottom else
  match e with | Const x  $\rightarrow$   $\beta$  x | Unknown  $\rightarrow$  top | Var v  $\rightarrow$  get m_0^\# v
  | BinOp op x y  $\rightarrow$  abstract_binop op (asem_expr m_0^\# x) (asem_expr m_0^\# y)

```

Backward analysis Our aim is to have an instance for our memory of mem_{adom} : it expects an `assume_` operator. Thus, below a backward analysis is defined for expressions. Given an expression e , an abstract value $r^\#$ and a memory $m_0^\#$, $\overleftarrow{\text{asem}} e r^\# m_0^\#$ computes a new abstract memory. That abstract memory refines the abstract values held in $m_0^\#$ as much as possible under the hypothesis that e lives in $r^\#$. The soundness of this analysis is encoded as a refinement on the output memory. Given any concrete memory m_0 and integer v approximated by $r^\#$, if the operational semantics of e at memory m_0 contains v , then m_0 should also be approximated by the output memory.

When e is a constant which is not contained in the concretization of the target abstract value $r^\#$, the hypothesis " e lives in $r^\#$ " is false, thus we translate that fact by outputting the unreachable memory `bottom`. In opposition, when e is `Unknown`, the hypothesis brings no new knowledge, thus we return the initial memory $m_0^\#$. In the case of a variable lookup (i.e. $e = \text{Var } v$ for some v), we consider $x^\#$, the abstract value living at v . Since our goal is to craft the most precise memory such that `Var v` is approximated by $r^\#$, we alter $m_0^\#$ by assigning $x^\# \sqcap r^\#$ at the variable v . Finally, in the case of binary operations, we make use of the backward operators and of recursion. Note that it is the only place where we need to insert a hint for the SMT solver: we assert an equality by asking F^* to normalize the terms. We state explicitly that the operational semantics of a binary operation reduces to two existentials: we manually unfold the definition of $\text{osem}_{\text{expr}}$ and `lift_binop`. The `decreases` clause explains to F^* why and how the recursion terminates.

```

let rec  $\overleftarrow{\text{asem}}$  {|\text{num}_{\text{adom}} \tau|\} (e: expr) (r^\#:  $\tau$ ) (m_0^\#: amem  $\tau$ )
: Tot (m_1^\#: amem  $\tau$  { (* decreases: *) m_1^\#  $\sqsubseteq$  m_0^\#  $\wedge$  (* soundness: *)
  ( $\forall$  (m_0: mem) (v: int_m). (v  $\in$   $\gamma$  r^\#  $\wedge$  m_0  $\in$   $\gamma$  m_0^\#  $\wedge$  v  $\in$   $\text{osem}_{\text{expr}} m_0 e$ )
   $\implies$  m_0  $\in$   $\gamma$  m_1^\#)}) (decreases e)
= if m_0^\#  $\sqsubseteq$  bottom then bottom else match e with
| Const x  $\rightarrow$  if cgamma r^\# x then m_0^\# else bottom | Unknown  $\rightarrow$  m_0^\#
| Var v  $\rightarrow$  let x^\#:  $\tau = r^\# \sqcap$  get m_0^\# v in
  if x^\#  $\sqsubseteq$  bottom then Bot else amem_update v x^\# m_0^\#

```



```

| BinOp op ex ey → let  $\overleftarrow{\text{op}}$  =  $\overleftarrow{\text{abstract\_binop}}$  op in
  let  $x^\#, y^\#$  =  $\overleftarrow{\text{op}}$  (asemexpr  $m_0^\#$  ex) (asemexpr  $m_0^\#$  ey)  $r^\#$  in
  let  $r^\#$ : amem  $\tau$  =  $\overleftarrow{\text{asem}}$  ex  $x^\#$   $m_0^\#$   $\sqcap$   $\overleftarrow{\text{asem}}$  ey  $y^\#$   $m_0^\#$  in
  assert_norm ( $\forall$  (m: mem) (v: intm). v ∈ osemexpr m e
     $\iff$  ( $\exists$  (x y: intm). x ∈ osemexpr m ex  $\wedge$  y ∈ osemexpr m ey
       $\wedge$  v == concrete_binop op x y));
   $r^\#$ 

```

Iterating the backward analysis While a concrete test is idempotent, it is not the case for abstract ones. Our goal is to refine an abstract memory under a hypothesis as much as possible. Since $\overleftarrow{\text{asem}}$ is proven sound and decreasing, we can repeat the analysis as much as we want. We introduce `prefixpoint` that computes a pre-fixpoint. However, even if the function from which we want to get a prefixpoint is decreasing, this is not a guarantee for termination. The type `measure` below associates an order to a measure that ensures termination. Such a measure cannot be implemented for a lattice that has infinite decreasing or increasing chains. We also require a maximum for this measure, so that we can reverse the measure easily in the context of postfixpoints iteration.

```

type measure #a (ord: a → a → bool)
= { f: f: (a → ℕ) { $\forall$  x y. x `ord` y  $\implies$  x  $\neq$  y  $\implies$  f x < f y}
  ; max: (max: ℕ { $\forall$  x. f x < max}) }

```

Let us focus on `prefixpoint`: given an order \sqsubseteq with its measure m , it iterates a decreasing function f , starting from a value x . The argument r is a binary relation which is required to hold for every couple $(x, f x)$. r is also required to be transitive, so that morally $r x (f^n x)$ holds. `prefixpoint` is specified to return a prefixpoint y , that is, with $r x y$ holding.

```

let rec prefixpoint (( $\sqsubseteq$ ): order  $\tau$ ) (m: measure ( $\sqsubseteq$ ))
  (r:  $\tau \rightarrow \tau \rightarrow \text{prop}$  {trans r}) (f:  $\tau \rightarrow \tau$  { $\forall e$ . f e  $\sqsubseteq$  e  $\wedge$  r e (f e)}) (x:  $\tau$ )
  : Tot (y:  $\tau$  {r x y  $\wedge$  f y == y  $\wedge$  y  $\sqsubseteq$  x}) (decreases (m.f x))
= let x' = f x in if x  $\sqsubseteq$  x' then x else prefixpoint ( $\sqsubseteq$ ) m r f x'

```

Below is defined $\overleftarrow{\text{asem_fp}}$ the iterated version of $\overleftarrow{\text{asem}}$. Besides using `prefixpoint`, the only thing required here is to spell out t , the relation we want to ensure.

```

let  $\overleftarrow{\text{asem\_fp}}$  { $|\text{num}_{\text{adom}} \tau|$ } (e: expr) (r:  $\tau$ ) ( $m_0^\#$ : amem  $\tau$ )
  : Tot ( $m_1^\#$ : amem  $\tau$  {( $\forall$  ( $m_0$ : mem) (v: intm).  $m_1^\# \sqsubseteq m_0^\# \wedge$ 
    ( $v \in \gamma r \wedge m_0 \in \gamma m_0^\# \wedge v \in \text{osem}_{\text{expr}} m_0 e$ )  $\implies m_0 \in \gamma m_1^\#$ )}))
= let t ( $m_0^\# m_1^\#$ : amem  $\tau$ ) =  $\forall$  (m: mem) (v: intm).
  ( $v \in \gamma r \wedge m \in \gamma m_0^\# \wedge v \in \text{osem}_{\text{expr}} m e$ )  $\implies m \in \gamma m_1^\#$  in
  prefixpoint corder order_measure t ( $\overleftarrow{\text{asem}}$  e r)  $m_0^\#$ 

```

A `memadom` instance We defined both a forward and backward analysis for expressions. Implementing an `memadom` instance for `amem` is thus easy, as shown

below. For any numerical abstract domain τ , `amemory_mem_adom` provides an `memadom`, that is, an abstract domain for memories, providing nontrivial proofs of correctness. Still, this is proven automatically.

```
instance amemory_mem_adom {! nd: numadom  $\tau$  !}: memadom (amem  $\tau$ ) =
  let adom: adom (amem  $\tau$ ) = amemadom in { maadom = adom; mamem = ()
  ; assume_ = ( $\lambda m^\# e \rightarrow \overleftarrow{\text{asem\_fp}} e \text{ gt}_0 m^\# \sqcup \overleftarrow{\text{asem\_fp}} e \text{ lt}_0 m^\#$ )
  ; assign = ( $\lambda m^\# v e \rightarrow \text{let } v^\#: \tau = \text{asem\_expr } m^\# e \text{ in}$ 
    if  $v^\# \sqsubseteq \text{bottom}$  then Bot else amem_update v  $v^\# m^\#$ )}
```

8 Statement Abstract Interpretation

Wrapping up the implementation of our abstract interpreter, this section presents the abstract interpretation of IMP statements. For every memory type μ that instantiates the typeclass of abstract memories `memadom`, the abstract semantics `asemstmt` maps statements and initial abstract memories to final memories. `memadom` is defined and proven correct below.

Given a statement `s`, and an initial abstract memory $m_0^\#$, `memadom s m0#` is a final abstract memory so that for any initial concrete memory `m` approximated by $m_0^\#$ and for any acceptable final concrete memory `m'` considering the operational semantics, `m'` is approximated by `memadom s m0#`. Here, we give two hints to the SMT solver: by normalization (`assert_norm`), we unfold the operational semantics in the case of choices or sequences. The analysis of an assignment or an assume is very easy since we already have operators defined for these cases. In the case of the sequence of two statements, we simply recurse. Similarly, when the statement is a choice, we recurse on its two possibilities. Then the two resulting abstract memories are merged back together. The last case to be handled is the loop, that is some statement of the shape `Loop body`. We compute a fixpoint $m_1^\#$ for `body`, by widening: it therefore approximates correctly the operational semantics of `Loop body`, since it is defined as a transitive closure. F*'s standard library provides the lemma `stable_on_closure`; of which we give a simplified signature below. The concretization $\gamma m_1^\#$ is a set, that is a predicate: we use this lemma with $\gamma m_1^\#$ as predicate `p` and with the operational semantics as relation `r`.

```
val simplified_stable_on_closure: r:( $\tau \rightarrow \tau \rightarrow \text{prop}$ )  $\rightarrow$  p:( $\tau \rightarrow \text{prop}$ )
 $\rightarrow$  Lemma (requires  $\forall x y. p x \wedge r x y \implies p y$ )
      (ensures  $\forall x y. p x \wedge \text{closure } r x y \implies p y$ )
```

```
let rec asemstmt {! md: memadom  $\mu$  !} (s: stmt) (m0#:  $\mu$ )
: (m1#:  $\mu \{ \forall (m m': \text{mem}). (m \in \gamma m_0^\# \wedge m' \in \text{osem}_{\text{stmt}} s m) \implies m' \in \gamma m_1^\# \}$ )
= assert_norm ( $\forall s_0 s_1 (m_0 \text{ mf}: \text{mem}). \text{osem}_{\text{stmt}} (\text{Seq } s_0 s_1) m_0 \text{ mf}$ 
  == ( $\exists (m_1: \text{mem}). m_1 \in \text{osem}_{\text{stmt}} s_0 m_0 \wedge \text{mf} \in \text{osem}_{\text{stmt}} s_1 m_1$ ));
  assert_norm ( $\forall a b (m_0 \text{ mf}: \text{mem}). \text{osem}_{\text{stmt}} (\text{Choice } a b) m_0 \text{ mf}$ 
  == ( $\text{mf} \in (\text{osem}_{\text{stmt}} a m_0 \cup \text{osem}_{\text{stmt}} b m_0)$ ));
if m0#  $\sqsubseteq$  bottom then bottom
```

```

else match s with
| Assign v e → assign m0# v e
| Assume e → assume_ m0# e | Seq s t → asemstmt t (asemstmt s m0#)
| Choice a b → asemstmt a m0# ⊔ asemstmt b m0#
| Loop body → let m1#: μ = postfixpoint corder order_measure
                (λ(m#:μ) → widen m# (asemstmt body m# <: μ))
                in stable_on_closure (osemstmt body) (γ m1#) (); m1#

```

Below we show the definition of `postfixpoint`, which is similar to `prefixpoint`. However, it is simpler because it only ensures its outcome is a postfixpoint.

```

let rec postfixpoint ((⊑): order τ) (m: measure (⊑))
  (f: τ → τ {∀ x. x ⊑ f x}) (x: τ)
  : Tot (y: τ {f y == y ∧ (⊑) x y}) (decreases (m.max - m.f x))
  = let x' = f x in if x' ⊑ x then x else postfixpoint (⊑) m f x'

```

9 Conclusion and further works

We presented almost the entire code of our abstract interpreter for IMP. Our approach to abstract interpretation is concretization-based, and follows the methodology of [3,11]. While using F*, we did not encounter any issue regarding expressiveness, and additionally gained a lot in proof automatization, to finally implement a fairly modular abstract interpreter. The table below compares the line-of-proof vs. line-of-code ratio of our implementation compared to some of the available verified abstract interpreters. Ours is up to 17 times more proof efficient. It is very compact, and requires a negligible amount of manual proofs. This comparison has its limits, since the different formalizations do not target the same programming languages: [11] and [3] handles the full C language, while [5] and the current paper deal with more simple imperative languages. Also, proof effort usually does not scale linearly.

	Code	Proof	Ratio	Feature set
This paper	487	39	0.08	Simple imperative language
Pichardie et al. [5]	3 725	5 020	1.35	Simple imperative language
Verasco [11]	16 847	17 040	1.01	CompCert C langage
Blazy et al. [3]	4 000	3 500	0.87	CompCert C langage

The sources of our abstract interpreter sources are available along with a set of example programs; building it natively or as a web application is easy, reproducible² and automated.

This work is very far from the scope of Verasco which required about four years of human time [12,10], but our results, which required 3 months of work with F* expertise, are very encouraging.

² Our build process relies on the purely functional Nix package manager.

Further work We aim at following the path of Verasco by adding real-world features to our abstract interpreter and consider a more realistic target language such as one of the CompCert C-like input languages. One of the weaknesses of Verasco is its efficiency. Using Low*, a C DSL for F*, it is possible to write (with a nontrivial additional effort related to Low*) a very efficient C and formally verified abstract interpreter. This development also opens the path for enriching F* automation via verified abstract interpretation.

Acknowledgements

This work is supported by a European Research Council (ERC) Consolidator Grant for the project VESTA, funded under the European Union's Horizon 2020 Framework Programme (grant agreement 772568).

References

1. *Provably secure communication software*, <https://project-everest.github.io/>
2. *Supplementary materials*: <https://zenodo.org/record/5168401>
3. S. Blazy, V. Laporte, A. Maroneze, D. Pichardie: *Formal verification of a C value analysis based on abstract interpretation*. In: SAS. pp. 324–344. LNCS (2013)
4. B. Bond, C. Hawblitzel, M. Kapritsos, R. Leino, J. Lorch, B. Parno, A. Rane, S. Setty, L. Thompson: *Vale: Verifying high-performance cryptographic assembly code*. In: Proceedings of the USENIX Security Symposium. USENIX (August 2017), distinguished Paper Award
5. D. Cachera, D. Pichardie: *A certified denotational abstract interpreter*. In: Proc. of International Conference on Interactive Theorem Proving (ITP-10). LNCS, vol. 6172, pp. 9–24 (2010)
6. P. Cousot, R. Cousot, J. Feret, L. Mauborgne, A. Miné, D. Monniaux, X. Rival: *The astrée analyzer*. In: European Symposium on Programming. pp. 21–30 (2005)
7. D. Darais, M. Might, D. Van Horn: *Galois transformers and modular abstract interpreters: Reusable metatheory for program analysis*. In: Proceedings of the 2015 ACM SIGPLAN International Conference on Object-Oriented Programming, Systems, Languages, and Applications. pp. 552–571. OOPSLA 2015 (2015)
8. P. David: *Interprétation abstraite en logique intuitionniste : extraction d’analyseurs Java certifiés*. Ph.D. thesis, Université Rennes 1 (2005), in french
9. L. De Moura, N. Bjørner: *Z3: An efficient smt solver*. In: International conference on Tools and Algorithms for the Construction and Analysis of Systems. pp. 337–340 (2008)
10. J.H. Jourdan: *Verasco: a Formally Verified C Static Analyzer*. Theses, Université Paris Diderot-Paris VII (May 2016)
11. J.H. Jourdan, V. Laporte, S. Blazy, X. Leroy, D. Pichardie: *A formally-verified C static analyzer*. In: 42nd symposium Principles of Programming Languages. pp. 247–259. ACM Press (2015)
12. V. Laporte: *Verified static analyzers for low-level languages*. Theses, Université Rennes 1 (Nov 2015)
13. G. Martínez, D. Ahman, V. Dumitrescu, N. Giannarakis, C. Hawblitzel, C. Hritcu, M. Narasimhamurthy, Z. Paraskevopoulou, C. Pit-Claudel, J. Protzenko, T. Ramananandro, A. Rastogi, N. Swamy: *Meta-F*: Proof automation with SMT, tactics, and metaprograms*. In: 28th European Symposium on Programming (ESOP). pp. 30–59 (2019)
14. T. Nipkow: *Abstract interpretation of annotated commands*. In: Beringer, Felty (eds.) Interactive Theorem Proving (ITP 2012). vol. 7406, pp. 116–132 (2012)
15. J. Protzenko, B. Parno, A. Fromherz, C. Hawblitzel, M. Polubelova, K. Bhargavan, B. Beurdouche, J. Choi, A. Delignat-Lavaud, C. Fournet, N. Kulatova, T. Ramananandro, A. Rastogi, N. Swamy, C.M. Wintersteiger, S. Zanella-Béguelin: *Evercrypt: A fast, verified, cross-platform cryptographic provider*. In: IEEE Symposium on Security and Privacy. IEEE (May 2020)
16. J.K. Zinzindohoué, K. Bhargavan, J. Protzenko, B. Beurdouche: *Hacl: A verified modern cryptographic library*. In: Proceedings of the 2017 ACM SIGSAC Conference on Computer and Communications Security. p. 1789–1806. CCS ’17, Association for Computing Machinery (2017). <https://doi.org/10.1145/3133956.3134043>