

Monadic Abstract Interpreters

Ilya Sergey

IMDEA Software Institute, Spain
ilya.sergey@imdea.org

Dominique Devriese

iMinds – DistriNet, KU Leuven, Belgium
dominique.devriese@cs.kuleuven.be

Matthew Might

University of Utah, USA
might@cs.utah.edu

Jan Midtgaard

Aarhus University, Denmark
jmi@cs.au.dk

David Darais

Harvard University, USA
darais@seas.harvard.edu

Dave Clarke Frank Piessens

iMinds – DistriNet, KU Leuven, Belgium
{firstname.lastname}@cs.kuleuven.be

Abstract

Recent developments in the systematic construction of abstract interpreters hinted at the possibility of a broad unification of concepts in static analysis. We deliver that unification by showing context-sensitivity, polyvariance, flow-sensitivity, reachability-pruning, heap-cloning and cardinality-bounding to be independent of any particular semantics. Monads become the unifying agent between these concepts and between semantics. For instance, by plugging the same “context-insensitivity monad” into a monadically-parameterized semantics for Java or for the lambda calculus, it yields the expected context-insensitive analysis.

To achieve this unification, we develop a systematic method for transforming a concrete semantics into a monadically-parameterized abstract machine. Changing the monad changes the behavior of the machine. By changing the monad, we recover a spectrum of machines—from the original concrete semantics to a monovariant, flow- and context-insensitive static analysis with a singly-threaded heap and weak updates.

The monadic parameterization also suggests an abstraction over the ubiquitous monotone fixed-point computation found in static analysis. This abstraction makes it straightforward to instrument an analysis with high-level strategies for improving precision and performance, such as abstract garbage collection and widening.

While the paper itself runs the development for continuation-passing style, our generic implementation replays it for direct-style lambda-calculus and Featherweight Java to support generality.

Categories and Subject Descriptors F.3.2 [Logics and Meanings of Programs]: Semantics of Programming Languages—Program analysis, Operational semantics

General Terms Languages, Theory

Keywords abstract machines, abstract interpretation, monads, operational semantics, collecting semantics, abstract garbage collection, interpreters

1. Introduction

Recent work on systematizing the construction of abstract interpreters [15, 23, 24] hints at the possibility of broad theoretical unification within static analysis. Van Horn and Might [23] sketch a method for abstracting abstract machines into static analyzers by bounding the store of the abstract machine: once the store is bounded, the abstraction and then the analysis follows.

But, bounding the store is an act of design—of human intervention. How a designer bounds the store immediately determines classical properties of the analysis such as its context-sensitivity and its polyvariance. While not directly expressed in terms of a bound on the store, other classical properties are also related to the abstraction and handling of the store, including heap-cloning, reachability-pruning and cardinality-bounding. But, other properties, such as flow-sensitivity, path-sensitivity and some kinds of widening, have little to do with the store, and more to do with the (re-)interpretation of the abstract semantics during analysis.

Fortunately, there is a construct that encompasses all of these concerns: the monad. Monads were originally adapted to programming languages to provide a durable abstraction for mutation in a purely functional language. As such, expressing a semantics monadically allows the monad to fully veil the details of interacting with a store. Yet monads have always provided more than just a means for hiding effects in a purely functional manner: they also allow a near-complete reinterpretation of computations expressed monadically, *e.g.*, the instant and powerful non-determinism of the list monad. This semantic reflection in precisely the right dimensions allows monads to encapsulate a variety of concepts in static analysis.

The payoff of this realization is immediate: we can monadically refactor semantics for languages as diverse as the lambda calculus and Java, yet define notions like context-sensitivity for both at the same time, with the same monad.

Our presentation details every step of the monadic refactoring for continuation-passing style lambda calculus, and then develops the monadic parameters that induce static analyzers. The implementation of the approach in the accompanying code repository replays the same monadic refactoring for a direct-style lambda calculus and for Featherweight Java. *The monads remain the same.*

We have chosen Haskell in lieu of formal mathematics for the presentation for two reasons: (1) Haskell is directly executable, and (2) Haskell has concise, convenient and readable syntax for expressing monads—the central actor in our work. Our fundamental results are no more restricted to Haskell than monads are restricted to Haskell.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

PLDI'13, June 16–19, 2013, Seattle, WA, USA.

Copyright © 2013 ACM 978-1-4503-2014-6/13/06...\$10.00

1.1 Overview

Van Horn and Might’s systematic abstraction relies on breaking down recursive structure in the concrete state-space and threading that structure through a store. This refactoring of the state-space imposes a corresponding store-passing-style transformation on the semantics.

Were we writing an *interpreter in Haskell* rather than a *formal semantics*, we would recognize store-passing style as an anti-pattern: in Haskell, monads are the right generalization for capturing artifacts like a store in a purely functional manner. When we apply a monadic transformation in lieu of a store-passing transformation to a concrete semantics, the resulting monad erupts as the catalyst for unifying what appeared as *ad hoc* choices in the design of classical static analyses. Even the nondeterminism that arises during abstraction of an operational semantics can be captured, explained and throttled entirely monadically.

By delegating interaction with the store into a monad, the monad determines not only the polyvariance and context-sensitivity of the analysis, but also the degree to which the analysis prunes the abstract heap based on reachability (abstract garbage collection) and bounds the cardinality of abstract addresses for shape analysis. In fact, it suggests a refactoring of the traditional fixed-point iteration such that path-sensitivity and flow-sensitivity also fall out as natural parameters.

Our monadic abstraction of the semantics orthogonalizes the classic dimensions of static analysis, independent of the specific semantics in use.

We illustrate a systematic method for transforming a concrete semantics into a monadically-parameterized machine, comprising both *concrete* and *abstract* interpretation, such that the monad determines the classical properties of an analysis. As in recent work on abstracting abstract machines [23], our semantic transformation implicitly utilizes the techniques of Danvy *et al.* [1, 2, 7] in order to calculate an abstract¹ machine equivalent to the concrete semantics. It diverges with this line of work by applying a monadic-normal-form transformation [6] (instead of a store-passing-style transformation) to the rules for the machine.

1.2 Contributions

- The central theoretical contribution of the paper is identifying and employing monads as a mechanism to abstract over the fundamental characteristics of the analysis.
- The central practical contribution is an executable proof-of-concept implementation of the described decomposition for a series of calculi.²
- We decouple the interpretation of the semantics from a monotonic fixed-point computation, which makes it possible to define analysis widening strategies independently of the semantics and of the analysis’ other parameters.
- We illustrate degrees of freedom when constructing the analysis using our framework and show that components implementing non-deterministic transitions, polyvariance and abstract counting are semantics-independent and can be reused for different calculi (*e.g.*, Java and the lambda calculus) and analysis families.

¹Abstract in the sense that it models the salient intensional behavior of a program.

²The implementation is available for the lambda-calculus (both in the form of CPS [15] and CESK-machine [23]) and Featherweight Java [19]:

<http://github.com/ilyasergey/monadic-cfa>

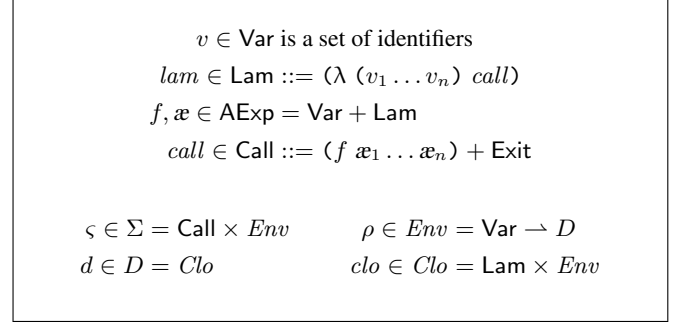


Figure 1: A grammar for CPS and a concrete state-space.

2. Setting the Scene: Analyzing Continuation-Passing Style

We start our discovery of the monadic refactoring process with a minimalist higher-order language: continuation-passing style λ -calculus (CPS). We will apply the systematic abstraction process as described by Van Horn and Might [23] in full to CPS. Afterward, we’ll modify the process by transforming the semantics into monadic normal form after the store-passing transformation. Thus, all interaction with the store will pass through a monad. Because interaction with the store is central to describing facets of modern flow analysis, such as context-sensitivity and heap-cloning, we will be able to describe these facets independently of a particular semantics.

In CPS (Figure 1), the lambda calculus is partitioned into two worlds: call sites and atomic expressions. Atomic expressions are lambda terms and variable references. Call sites encode the application of a function to arguments. A classical abstract machine for CPS [9] needs only two components in its state-space Σ —see Figure 1—a control component (Call) and an environment (Env). The domain D contains denotable values. CPS is so simple that there is only one kind of denotable value—closures.

The injector $\mathcal{I} : \text{Call} \rightarrow \Sigma$ maps a program into this state-space:

$$\mathcal{I}(\text{call}) = (\text{call}, []),$$

In CPS, there is only one rule to describe the transition relation, $(\Rightarrow) \subseteq \Sigma \times \Sigma$. We don’t write index ranges explicitly for series of meta-expressions, assuming $i = 1..n$ is obvious from the context.

$$\begin{aligned}
 (\llbracket (f \mathfrak{x}_1 \dots \mathfrak{x}_n) \rrbracket, \rho) & \Rightarrow (\text{call}, \rho''), \text{ where} \\
 (\llbracket (\lambda (v_1 \dots v_n) \text{ call}) \rrbracket, \rho') & = \mathcal{A}(f, \rho) \\
 d_i = \mathcal{A}(\mathfrak{x}_i, \rho) & \quad \rho'' = \rho' [v_i \mapsto d_i],
 \end{aligned}$$

where the evaluator $\mathcal{A} : \text{AExp} \times \text{Env} \rightarrow \text{Clo}$ evaluates an atomic expression:

$$\mathcal{A}(v, \rho) = \rho(v) \quad \mathcal{A}(\text{lam}, \rho) = (\text{lam}, \rho).$$

2.1 Attempting structural abstraction

A structural abstraction carries abstraction component-wise across the state-space and then lifts naturally over internal domains. However, a structural abstraction of the state-space for CPS yields:

$$\begin{aligned}
 \hat{\Sigma} & = \text{Call} \times \widehat{\text{Env}} & \widehat{\text{Env}} & = \text{Var} \rightarrow \hat{D} \\
 \hat{D} & = \mathcal{P}(\widehat{\text{Clo}}) & \widehat{\text{Clo}} & = \text{Lam} \times \widehat{\text{Env}}.
 \end{aligned}$$

A structural abstraction preserves mutual recursion between closures and environments, and with it, the unboundedness of the abstract state-space. Since our goal with this abstraction was a finite abstract state-space (hence a trivially computable abstract semantics), structural abstraction alone is insufficient.

2.2 Cutting recursion with a store

To arrive at a finite state-space, the systematic method detailed in “Abstracting Abstract Machines” [23] calls for transforming the abstract machine into store-passing style, which cuts the recursive knot between environments and closures by introducing addresses. With the introduction of a store, $\sigma \in Store = Addr \rightarrow D$, the store-passing transform produces the following state-space:

$$\begin{aligned} \zeta \in \Sigma &= Call \times Env \times Store \\ \rho \in Env &= Var \rightarrow Addr \\ d \in D &= Clo \\ clo \in Clo &= Lam \times Env \\ a \in Addr &\text{ is an infinite set of addresses.} \end{aligned}$$

The modification of the transition relation is straightforward:

$$\begin{aligned} \overbrace{(\llbracket (f \ x_1 \dots x_n) \rrbracket, \rho, \sigma)}^{\zeta} &\Rightarrow (call, \rho'', \sigma'), \text{ where} \\ (\llbracket (\lambda (v_1 \dots v_n) \ call) \rrbracket, \rho') &= \mathcal{A}(f, \rho, \sigma) \\ d_i &= \mathcal{A}(x_i, \rho, \sigma) \quad a_i = alloc(v_i, \zeta) \\ \rho'' &= \rho'[v_i \mapsto a_i] \quad \sigma' = \sigma[a_i \mapsto d_i]. \end{aligned}$$

The evaluator $\mathcal{A} : AExp \times Env \times Store \rightarrow D$ is modified to take the store as an additional argument:

$$\mathcal{A}(v, \rho, \sigma) = \sigma(\rho(v)) \quad \mathcal{A}(lam, \rho, \sigma) = (lam, \rho).$$

And, the (presently opaque) address-allocator $alloc : Var \times \Sigma \rightarrow Addr$ yields a fresh address for each variable.

2.3 A second attempt at abstraction

With the recursion sliced from the state-space by the store-passing transformation, a structural abstraction succeeds in producing a finite abstract state-space:

$$\begin{aligned} \hat{\zeta} \in \hat{\Sigma} &= Call \times \widehat{Env} \times \widehat{Store} \\ \hat{\rho} \in \widehat{Env} &= Var \rightarrow \widehat{Addr} \\ \hat{\sigma} \in \widehat{Store} &= \widehat{Addr} \rightarrow \mathcal{P}(\widehat{D}) \\ \hat{d} \in \widehat{D} &= \widehat{Clo} \\ \widehat{clo} \in \widehat{Clo} &= Lam \times \widehat{Env} \\ \hat{a} \in \widehat{Addr} &\text{ is a finite set of abstract addresses,} \end{aligned}$$

and it induces a straightforward abstract transition relation:

$$\begin{aligned} \overbrace{(\llbracket (f \ x_1 \dots x_n) \rrbracket, \hat{\rho}, \hat{\sigma})}^{\hat{\zeta}} &\rightsquigarrow (call, \hat{\rho}'', \hat{\sigma}'), \text{ if} \\ (\llbracket (\lambda (v_1 \dots v_n) \ call) \rrbracket, \hat{\rho}') &\in \hat{\mathcal{A}}(f, \hat{\rho}, \hat{\sigma}) \\ \hat{d}_i \in \hat{\mathcal{A}}(x_i, \hat{\rho}, \hat{\sigma}) &\quad \hat{a}_i = \widehat{alloc}(v_i, \hat{\zeta}) \\ \hat{\rho}'' &= \hat{\rho}'[v_i \mapsto \hat{a}_i] \quad \hat{\sigma}' = \hat{\sigma} \sqcup [\hat{a}_i \mapsto \{\hat{d}_i\}]. \end{aligned}$$

Branching to every possible abstract closure introduces a subtle nondeterminism.

Naturally, the abstract argument evaluator wraps closures as singletons, but looks up variables as in the concrete version:

$$\hat{\mathcal{A}}(v, \hat{\rho}, \hat{\sigma}) = \hat{\sigma}(\hat{\rho}(v)) \quad \hat{\mathcal{A}}(lam, \hat{\rho}, \hat{\sigma}) = \{\{lam, \hat{\rho}\}\}.$$

The join on the abstract store allows each address in the finite set of abstract addresses to soundly represent multiple concrete addresses:

$$\hat{\sigma} \sqcup \hat{\sigma}' = \lambda \hat{a}. \hat{\sigma}(\hat{a}) \cup \hat{\sigma}'(\hat{a}).$$

And, it should now be clear that the abstract allocator $\widehat{alloc} : Var \times \hat{\Sigma} \rightarrow \widehat{Addr}$ determines the polyvariance of the analysis (*i.e.*, distinguishing between bindings of the same variable in different evaluation contexts [22, 19]) because \widehat{alloc} determines how many abstract variants are associated with each variable. For instance, allocating each time a new address corresponds to associating just one value with a variable at each moment of the program execution.

2.3.1 Example: Monovariant analysis (OCFA)

For example, a classical monovariant analysis (OCFA) comes from defining the set of abstract addresses to be the set of variables ($\widehat{Addr}_{OCFA} = Var$), and then using variables as their own abstract addresses:

$$\widehat{alloc}_{OCFA}(v, \hat{\zeta}) = v.$$

2.4 Introducing context with time-stamps

Van Horn and Might [23] realize that the abstract state-space as it stands lacks sufficient information to instantiate classical strategies for polyvariance (like k -CFA). To fix this, they introduce time-stamps (to the concrete and abstract semantics) as a new component of the state.

Time-stamps are used to remember execution context, and are directly responsible for the context-sensitivity of the analysis:

$$\hat{\zeta} \in \hat{\Sigma} = Call \times \widehat{Env} \times \widehat{Store} \times \widehat{Time}$$

$\hat{t} \in \widehat{Time}$ is a finite set of abstract times.

Each transition augments the time through an opaque function, $\widehat{tick} : \widehat{Clo} \times \hat{\Sigma} \rightarrow \widehat{Time}$ and by giving the allocator $\widehat{alloc} : Var \times \widehat{Time} \rightarrow \widehat{Addr}$ access to this context instead of the whole state:

$$\begin{aligned} \overbrace{(\llbracket (f \ x_1 \dots x_n) \rrbracket, \hat{\rho}, \hat{\sigma}, \hat{t})}^{\hat{\zeta}} &\rightsquigarrow (call, \hat{\rho}'', \hat{\sigma}', \hat{t}'), \text{ if} \\ (\llbracket (\lambda (v_1 \dots v_n) \ call) \rrbracket, \hat{\rho}') &\in \hat{\mathcal{A}}(f, \hat{\rho}, \hat{\sigma}) \\ \hat{d}_i \in \hat{\mathcal{A}}(x_i, \hat{\rho}, \hat{\sigma}) &\quad \hat{t}' = \widehat{tick}(\widehat{clo}, \hat{\zeta}) \\ \hat{a}_i = \widehat{alloc}(v_i, \hat{t}') &\quad \hat{\rho}'' = \hat{\rho}'[v_i \mapsto \hat{a}_i] \\ \hat{\sigma}' &= \hat{\sigma} \sqcup [\hat{a}_i \mapsto \{\hat{d}_i\}]. \end{aligned}$$

2.4.1 Example: k -CFA-style context-sensitivity

The introduction of this component makes it possible to model k -CFA [22] by defining times to be sequences of up to k call sites and addresses to be pairs of variables and call sites:

$$\begin{aligned} \widehat{Time}_{k\text{-CFA}} &= Call^{\leq k} \\ \widehat{Addr}_{k\text{-CFA}} &= Var \times \widehat{Time}_{k\text{-CFA}}, \end{aligned}$$

and using corresponding definitions

$$\begin{aligned} \widehat{tick}_{k\text{-CFA}}(\widehat{clo}, (call, \dots, \hat{t})) &= \lfloor call : \hat{t} \rfloor_k \\ \widehat{alloc}_{k\text{-CFA}}(v, (\dots, \hat{t}')) &= (v, \hat{t}') \end{aligned}$$

where $\lfloor \cdot \rfloor_k$ limits its argument to at most length k .

3. Abstracting through a Monad

In flow analysis, interaction with the store determines essential properties of the analysis. Thus, by abstracting away interaction with the store through a monad, we introduce an abstraction layer for these essential properties.

In pure functional programming, the introduction of store-passing style to mimic side effects is an anti-pattern. The established remedy for this anti-pattern is the use of monadic form in

conjunction with a state-transformer monad. We can apply this remedy to our semantics as well.

To make the presentation unambiguous (and executable), we first transliterate the abstract semantics (for k -CFA at the moment) into Haskell, starting with the syntax for CPS:

```

type Var      = String
data Lambda = [Var] => CExp deriving (Eq, Ord)
data AExp    = Ref Var
              | Lam Lambda deriving (Eq, Ord)
data CExp    = Call AExp [AExp]
              | Exit      deriving (Eq, Ord)

```

The set `Call` from the definition of the state-space corresponds to the type `CExp` in the implementation. The type for the state-space is a 4-tuple:

```

type Σ      = (CExp, Env, Store, Time)
type k → v  = Map k v
type Env    = Var → Addr
type Store  = Addr → P Val
data Val    = Clo (Lambda, Env)
              deriving (Eq, Ord)
type Addr  = (Var, Time)
type Time  = [CExp]

```

where we utilize Haskell's unicode support to keep the implementation and the math as close as possible (e.g., writing `P Val` instead of `Set Val`). The transition relation becomes a function `next` on states into lists of states:

```

next :: Σ → [Σ]
next ς@(Call f aes, ρ, σ, t) = [(call, ρ'', σ', t') |
  proc@(Clo (vs => call, ρ')) ← Set.toList (arg (f, ρ, σ)),
  let t' = tick (proc, ς)
      as = [alloc (v, t', proc, ς) | v ← vs]
      ds = [arg (æ, ρ, σ) | æ ← aes]
      ρ'' = ρ' // [v => a | v ← vs | a ← as]
      σ' = σ □ [a => d | a ← as | d ← ds]]
next ς = [ς]

```

The function `arg` is the transliteration of the argument evaluator, \hat{A} . The function `tick` is the transliteration of the \widehat{tick} function. And, the function `alloc` is the transliteration of the \widehat{alloc} function. Finally, the utility infix operator $(//) :: Ord\ k \Rightarrow (k \rightarrow v) \rightarrow [(k, v)] \rightarrow (k \rightarrow v)$ is used to update a map with a list of key-value pairs, and (\implies) is just a synonym for pair constructor.

3.1 Capturing nondeterminism in the monad

The function `next` uses the list comprehension notation to more closely mimic the formalism. The list comprehension notation is syntactic sugar for the list comprehension monad.

We can tiptoe into monadic normal form by expanding the list comprehension $[(call, \rho'', \sigma', t') | \dots]$ into its monadic form:

```

mnext :: Σ → [Σ]
mnext ς@(Call f aes, ρ, σ, t) = do
  proc@(Clo (vs => call, ρ')) ← Set.toList (arg (f, ρ, σ))
  let t' = tick (proc, ς)
      as = [alloc (v, t', proc, ς) | v ← vs]
      ds = [arg (æ, ρ, σ) | æ ← aes]
      ρ'' = ρ' // [v => a | v ← vs | a ← as]
      σ' = σ □ [a => d | a ← as | d ← ds]
  return (call, ρ'', σ', t')
mnext ς = return ς

```

Since function evaluation is the only source of nondeterminism, we can clean up these semantics by creating a new function `fun` as a special version of `arg` for evaluating functions:

```

fun :: (AExp, Env, Store) → [Val]
fun = Set.toList ◦ arg

```

Thus, we can rewrite `mnext`:

```

mnext :: Σ → [Σ]
mnext ς@(Call f aes, ρ, σ, t) = do
  proc@(Clo (vs => call, ρ')) ← fun (f, ρ, σ)
  ... -- the rest is unchanged

```

To reformulate this function fully into monadic normal form, we must refactor and recurry `fun`, `arg`, `tick`, `alloc` to return singleton lists as well (hence `fun` and `arg` now have the same type):

```

fun  :: (Env, Store) → AExp → [Val]
arg  :: (Env, Store) → AExp → [Val]
tick :: Val → State → [Time]
alloc :: (Time, Val, State) → Var → [Addr]

```

which allows us to rewrite `mnext` in monadic normal form, replacing list comprehensions with a standard function `mapM :: Monad\ m => (a → m\ b) → [a] → m\ [b]` from a monadic toolset.

```

mnext :: Σ → [Σ]
mnext ς@(Call f aes, ρ, σ, t) = do
  proc@(Clo (vs => call, ρ')) ← fun (ρ, σ) f
  t' ← tick proc ς
  let as = mapM (alloc (t', proc, ς)) vs
      ds = mapM (arg (ρ, σ)) aes
  ... -- the rest is unchanged

```

3.2 Pulling the store into the monad

In the previous section, we refactored the abstract semantics so that nondeterminism is explicitly threaded through the list monad. In this section, we subtly reformulate the semantics in terms of a `CPSInterface` type class that hides interaction with the store, abstracting over five functions that form a *semantic interface* of the CPS calculus:

```

class Monad\ m => CPSInterface\ m where
  fun  :: Env → AExp → m\ Val
  arg  :: Env → AExp → m\ Val
  (↦)  :: Addr → Val → m\ ()
  alloc :: Time → Var → m\ Addr
  tick :: Val → PΣ → m\ Time

```

where the type `PΣ` encodes a partial state (with no store):

$$P\Sigma = (CExp, Env, Time)$$

Under this interface, we can separate the specification of the abstract semantics from the specific details of how it interacts with the store:

```

mnext :: (CPSInterface\ m) => PΣ → m\ PΣ
mnext ps@(Call f aes, ρ, t) = do
  proc@(Clo (vs => call, ρ')) ← fun ρ f
  t' ← tick proc ps
  as ← mapM (alloc t') vs
  ds ← mapM (arg ρ) aes
  let ρ'' = ρ' // [v => a | v ← vs | a ← as]
      sequence [a ↦ d | a ← as | d ← ds]
  return (call, ρ'', t')
mnext ς = return ς

```

3.3 Pulling time into the monad

Like the store, time-stamps were another *ad hoc* addition to the original semantics to engineer them into a form favorable for static analysis. We can lift time-stamps out of partial states and into monads as well, so that $P\Sigma = (CExp, Env)$.

The analysis monad can now assume internal access to the current time, which simplifies the interface:

```
class Monad m => CPSInterface m a where
  fun  :: Env a -> AExp -> m Val
  arg  :: Env a -> AExp -> m Val
  (map) :: Addr -> Val -> m ()
  alloc :: Var -> m Addr
  tick  :: Val -> P\Sigma -> m ()
```

Because the allocator *alloc* can assume access to time (and the store) inside the monad, it no longer needs to take it as a parameter, leaving the variable to be bound as its *only remaining parameter*. The simplification in the analysis monad is reflected in a simplification of the definition for transition function *mnext* as well, as time is no longer a component of the abstract states datatype $P\Sigma$:

```
mnext :: (CPSInterface m) => P\Sigma -> m P\Sigma
mnext ps@(Call f aes, \rho) = do
  proc@(Clo (vs => call, \rho')) <- fun \rho f
  tick proc ps
  as <- mapM alloc vs
  ... -- the rest is unchanged
```

3.4 Abstracting over addresses

At this point, we have a monadically abstracted abstracted³ abstract machine. Plugging in a monad controls nondeterminism, context and access to the store. However, we still have one component left to be abstracted.

Until now, we have assumed that abstract addresses are represented by a fixed datatype. In practice, the nature of the addresses determines the polyvariance and context-sensitivity of the analysis, since they usually divide bindings according to different contexts (also known as *contours* [22]). For example, Shivers' ICFA allocates distinct contexts for each call site. Lakhotia *et al.* [12] introduce ℓ -contexts to build a static analysis for obfuscated x86 binaries, employing finite sequences of unique enclosed function calls. Finally, one can take a bounded set of naturals $\{n \in \mathbb{N} \mid n \leq N\}$ for some N as contexts, which will give a good precision for sufficiently big N .

Thus, we must abstract over the structure of addresses in order to cover at least all the options from above. Unfortunately, addresses form a part of the state-space as a codomain of *Env*; thus, we need to parameterize our semantic domains by the type of addresses—*a*:

```
type P\Sigma a = (CExp, Env a)
type Env a = Var -> a
data Val a = Clo (Lambda, Env a)
           deriving (Eq, Ord)
type Store a = a -> P (Val a)
```

With this shift, we are no longer attached to the *k*-CFA-like representation of addresses, as both *Addr* and *Time* are gone.

Of course, the signature of the *CPSInterface* monad and *mnext* change accordingly, although the body of *mnext* remains unchanged. The final definition of the *CPSInterface* class and the function *mnext* are represented in Figure 2 and are not going to change in the remainder of our story. What *is* going to change is the

```
class Monad m => CPSInterface m a where
  fun  :: Env a -> AExp -> m (Val a)
  arg  :: Env a -> AExp -> m (Val a)
  (map) :: a -> Val a -> m ()
  alloc :: Var -> m a
  tick  :: Val a -> P\Sigma a -> m ()

mnext :: CPSInterface m a => P\Sigma a -> m (P\Sigma a)
mnext ps@(Call f aes, \rho) = do
  proc@(Clo (vs => call', \rho')) <- fun \rho f
  tick proc ps
  as <- mapM alloc vs
  ds <- mapM (arg \rho) aes
  let \rho'' = \rho' // [v => a | v <- vs | a <- as]
  sequence [a map d | a <- as | d <- ds]
  return (call', \rho'')
mnext \vars = return \vars
```

Figure 2: Semantic interface and a small-step semantics of CPS in a monadic form.

implementation of the semantic interface *CPSInterface* as well as the choice of the underlying monad for the analysis logic.

4. Recovering a Concrete Interpreter

With non-determinism, interaction with the store and time pulled into an analysis monad, and addresses abstracted, our semantics has become highly parameterized. As a result, we can easily recover a concrete interpreter as a “sanity” check for adequacy. With the standard *IO* monad as an underlying analysis monad, we can use the “real” heap as the store and implement mutable references as a datatype *IOAddr* containing essentially a pointer:

```
data IOAddr = IOAddr {lookup :: IORef (Val IOAddr)}
```

We will also need two simple administrative functions in order to write to and read from *IOAddr*s:

```
readIOAddr :: IOAddr -> IO (Val IOAddr)
readIOAddr = readIORef o lookup

writeIOAddr :: IOAddr -> Val IOAddr -> IO ()
writeIOAddr = writeIORef o lookup
```

The implementation of the *CPSInterface* type class directly mimics the concrete semantics defined in Section 2.2:

```
instance CPSInterface IO IOAddr where
  fun \rho (Lam l) = return $ Clo (l, \rho)
  fun \rho (Ref v) = readIOAddr (\rho! v)

  arg \rho (Lam l) = return $ Clo (l, \rho)
  arg \rho (Ref v) = readIOAddr (\rho! v)

  addr map v = writeIOAddr addr v
  alloc v = liftM IOAddr $ newIORef \perp
  tick _ _ = return ()
```

When a new address is allocated, we pass the function *newIORef* \perp as the initial value of the new reference, since the actual value to be bound at this address is not yet defined at the moment of allocation. The function *tick* is a no-op: in the real world, time advances without our help.

With this interpretation of the semantic interface, the concrete interpreter is simply a recursively defined driver loop that iterates the semantic transition function *mnext* until an *Exit* state is reached.

³This is not a typo.

```

interpret :: CExp → IO (PΣ IOAddr)
interpret e = go (e, Map.empty)
  where go :: (PΣ IOAddr) → IO (PΣ IOAddr)
        go s = do s' ← mnext s
                 case s' of x@(Exit, _) → return x
                          y           → go y

```

5. Recovering a Collecting Semantics

In this section, we recall key notions from lattice theory and abstract interpretation, crucial for constructing an interpreter for concrete or abstract small-step collecting semantics of a program. We proceed by translating the theory into programs in Haskell and demonstrate an implementation of a simple collecting semantics.

5.1 Basics of lattice theory and abstract interpretation

A *complete lattice* $\langle C; \sqsubseteq, \perp, \top, \sqcup, \sqcap \rangle$ is a partial order $\langle C; \sqsubseteq \rangle$ such that there exists a least upper bound (or join) $\sqcup S$ and a greatest lower bound (or meet) $\sqcap S$ of all subsets $S \subseteq C$. In particular $\sqcup C = \top$ and $\sqcap C = \perp$.

A point x is a fixed point of a function f if $f(x) = x$. Given two partial orders, $\langle C, \sqsubseteq \rangle$ and $\langle A, \leq \rangle$, a function f of type $C \rightarrow A$ is monotone if $\forall x, y : x \sqsubseteq y \implies f(x) \leq f(y)$. By the Knaster-Tarski fixed-point theorem a monotone functional f over a complete lattice has a *least fixed point* $\text{lfp}_{\sqsubseteq} f = \sqcap \{x \mid f(x) \sqsubseteq x\}$. Algorithmically the least fixed point of a monotone function f over a complete lattice of finite height can be computed by Kleene iteration: $\perp \sqsubseteq f(\perp) \sqsubseteq f^2(\perp) \sqsubseteq f^3(\perp) \sqsubseteq \dots$ since

$$\text{lfp}_{\sqsubseteq} f = \bigsqcup_{i \geq 0} f^i(\perp). \quad (1)$$

A *Galois connection* is a pair of functions α, γ connecting two partial orders $\langle C, \sqsubseteq \rangle$ and $\langle A, \leq \rangle$, such that $\forall a, c : \alpha(c) \leq a \iff c \sqsubseteq \gamma(a)$. We typeset Galois connections as: $\langle C, \sqsubseteq \rangle \xleftrightarrow[\alpha]{\gamma} \langle A, \leq \rangle$.

Given a Galois connection $\langle C, \sqsubseteq \rangle \xleftrightarrow[\alpha]{\gamma} \langle A, \leq \rangle$ and a monotone function $F_c : C \rightarrow C$, an *abstract function* F_a can be constructed as $\alpha \circ F_c \circ \gamma \leq F_a$. Therefore, by the *fixed-point transfer theorem* [5], we have $\alpha(\text{lfp } F_c) \leq \text{lfp } F_a$ when F_a is monotone.

The reachable states collecting semantics for a set of program states Σ , a transition function $(\sim) \subseteq \Sigma \times \Sigma$, and a set of initial states $\Sigma_0 \subseteq \Sigma$ is defined on a complete lattice $\langle \mathcal{P}(\Sigma), \subseteq, \emptyset, \Sigma, \cup, \cap \rangle$ as a least fixed point of a monotone functional \mathcal{F} , where

$$\mathcal{F}(\mathcal{X}) = \Sigma_0 \cup \{s' \mid s \in \mathcal{X} \wedge s \sim s'\}. \quad (2)$$

\mathcal{F} is uncomputable in general, which makes it a usual starting point for construction of an abstract, monotone transition function $\widehat{\mathcal{F}}$, such that $\alpha \circ \mathcal{F} \circ \gamma \leq \widehat{\mathcal{F}}$ and computing an approximate fixed point as $\text{lfp } \widehat{\mathcal{F}}$.

In the remainder of the section we will systematically abstract over computation of a small-step collecting semantics for CPS.

5.2 Implementing lattices and fixed point computations

We define a type class for a lattice following its algebraic definition:

```

class Lattice a where
  ⊥  :: a
  ⊤  :: a
  (⊆) :: a → a → Bool
  (⊔) :: a → a → a
  (⊔) :: a → a → a

```

We provide the following instances of *Lattice* for the standard container types that we use heavily for systematic abstraction of abstract machines [15]: unit, pairs, powersets and maps:

```

instance Lattice ()
instance (Lattice a, Lattice b) => Lattice (a, b)
instance (Ord s, Eq s) => Lattice (P s)
instance (Ord k, Lattice v) => Lattice (k → v)

```

The monadic interpreter from Figure 2, which we presented at the end of Section 3, is reminiscent of the definition (2) of the collecting semantics. However, the fact it uses a monad and is not an endo-function, prevents us from using it directly as a relation (\sim) in a functional \mathcal{F} . What we can do is to explicitly make a separation of concerns between *mnext* as an implementation of an abstract transition function and a computation of a least fixed point. We define the type class *Collecting* to let the implementer of the analysis define the logic of initiating the semantics (*i.e.*, providing an initial set Σ_0 from (2)) and “making a step”:

```

class Collecting m a fp | fp → a, fp → m where
  applyStep :: (a → m a) → fp → fp
  inject :: a → fp

```

The functional dependencies $fp \rightarrow a$ and $fp \rightarrow m$ state that the choice of a domain fp determines in which monad m the step function should be interpreted [11], as well as what the transition’s domain and co-domain a should be.⁴ The computation of the collecting semantics as a least fixed point can be then defined directly from the Kleene iteration (1):

```

kleeneIt :: (Lattice a) => (a → a) → a
kleeneIt f = loop ⊥
  where loop c = let c' = f c in
                if c' ⊆ c then c else loop c'

```

Finally, the collecting semantics (2) translates gracefully to the following function, mapping a transition function *step* and an initial state c to a result of the analysis.

```

exploreFP :: (Lattice fp, Collecting m a fp) =>
  (a → m a) → a → fp
exploreFP step c = kleeneIt F
  where F s = inject c ⊔ applyStep step s

```

It is now straightforward to implement a function that, given an expression, runs the analysis:

```

runAnalysis :: (CPSInterface m a, Lattice fp,
  Collecting m (PΣ a) fp) =>
  CExp → fp
runAnalysis e = exploreFP mnext (e, Map.empty)

```

The signature of *runAnalysis* outlines precisely the three degrees of freedom that can be changed in order to obtain different collecting semantics:

1. A monad, accounting for nondeterminism and passing analysis-specific state components (*i.e.*, time and store),
2. An implementation of the semantic interface of a language (*e.g.*, the one in Figure 2), and
3. The analysis lattice and an implementation of a fixed point computation that extracts the result of a single step from the analysis monad and augments the lattice argument appropriately.

5.3 StorePassing—a simple implementation of a collecting semantics for CPS

At this point we are ready to reconstruct a simple abstract interpreter computing a collecting, store-passing semantics of CPS by

⁴ Alternatively, we could use associated types to express the same sort of dependencies [3].

gradually instantiating three main aspects of the analysis, outlined at the end of the previous section.

5.3.1 A two-level analysis monad

We start from constructing a monad for a simple collecting, store-passing semantics by employing standard monad transformers *StateT* and the list monad:

```
type StorePassing s g = StateT g (StateT s [])
```

The analysis parameters of type *s* and *g* carry state components. They carry the store and the analysis’ *guts* respectively, where the latter can contain for example a “time” value.

In a desugared form, the type *StorePassing* is equivalent to the functional type $g \rightarrow s \rightarrow [(a, g), s]$ for some *a*, *g* and *s*, so the stack representation of a monad should be read “inside-out”. That is, a value of type *StorePassing* produces a set (represented by a list) of results of type *a*, coupled with components of type *g* and *s*.

5.3.2 A simple implementation of a CPS semantic interface

To turn the *StorePassing* monad into an interpretation of the semantics, we have to implement the *CPSInterface* type class. So far, we choose to implement addresses as Haskell *Integers* for simplicity. The store is represented as a map from integer addresses to sets of values. The implementation of the semantic interface is provided below.

```
instance CPSInterface
  (StorePassing (Store Integer) Integer) Integer
  where
    fun ρ (Lam l) = return $ Clo (l, ρ)
    fun ρ (Ref v) = lift $ getsNDSet $ λσ → σ ! (ρ ! v)
    arg ρ (Lam l) = return $ Clo (l, ρ)
    arg ρ (Ref v) = lift $ getsNDSet $ λσ → σ ! (ρ ! v)
    a ↦ d        = lift $ modify $
      Map.insert a (singleton d)
    alloc v      = gets id
    tick proc ps = modify $ λt → t + 1
```

The function *getsNDSet* is a crux of handling non-determinism in a monadic implementation of the analysis as a stateful monad. It allows one to examine the state, get multiple results and non-deterministically choose one. Its signature needs explanation:

```
getsNDSet :: (MonadPlus m, MonadState s m) =>
  (s → Set a) → m a
```

However, the type class constraints before \Rightarrow provide a clue. One can see that *m* is required to be a state monad, *i.e.*, carry an implicit state component of type *s*, which can be accessed and modified. The argument type $s \rightarrow \text{Set } a$ accounts for the non-deterministic result of type *a* that might come out of examining the state *s*. Finally, *MonadPlus m* is a constraint that ensures that *m* has a notion of non-deterministic choice, so the obtained results can be combined in a wrapped result *m a*. What comes as a nice surprise is that *StorePassing* is an instance of both *MonadPlus* and *MonadState*.

lift is a standard Haskell function for explicit management of the monadic stack, which is used explicitly to disambiguate the targets of accesses to the monad stack [13, 21]. In our implementation of *CPSInterface* the outermost state is reserved to carry “guts” (*i.e.*, the time component), so we need to employ *lift* to access the store σ , located on a “second level” of the monad stack. Finally, *gets* and *modify* are standard higher-order functions that allow one to examine and modify internal state of the monad.

```
gets    :: MonadState s m => (s → a) → m a
modify :: MonadState s m => (s → s) → m ()
```

For instance, the implementation of *tick* modifies the time component on the first level of the monadic stack (therefore, no explicit lifting is required).⁵

5.3.3 Computing a collecting semantics of CPS

The last missing ingredient we need to compute the collecting semantics of CPS is the definition of a fixed point computation that uses the *StorePassing* monad. We can reach a fix-point of type $\mathcal{P} ((P\Sigma a, g), s)$ step by step with the following definitions of *applyStep* and *inject*.

```
instance (Ord s, Ord a, Ord g, HasInitial g, Lattice s) =>
  Collecting (StorePassing s g)
  (P\Sigma a)
  (P ((P\Sigma a, g), s)) where
  inject p = singleton $ ((p, initial), ⊥)
  applyStep step fp = joinWith runStep fp where
    runStep ((c, t), s) =
      Set.fromList $ runStateT (runStateT (step c) t) s
```

The implementation of *inject* instruments a provided state with initial “guts”, defined by the value *initial* of the class *HasInitial* for *g* and the lattice minimum \perp for *s*, and wraps it into a singleton set. The class *HasInitial* is defined as

```
class HasInitial g where initial :: g
```

and its implementation for *Integers* is trivial (*e.g.*, *initial* = 0).

The most interesting element of the implementation is the utility function *joinWith*, defined as follows:

```
joinWith :: (Lattice a) =>
  (b → a) → Set b → a
joinWith f = Set.foldr ((⊔) ∘ f) ⊥
```

That is, given a set of values of type *b* and a function *f* of type $b \rightarrow a$ for a lattice *a*, *joinWith* traverses the structure, applying *f* to each of its leaf elements and combines the results using the lattice join (\perp).

In the definition of *applyStep*, *joinWith* takes a function that simply passes the state ς and the components *t* and *s* to the provided function *step*, runs a monad and collects the result into a set. This function is applied to *all* states in *fp* and the results are joined.

All ingredients to run the analysis are now in place, and all we need to do is to use the function *runAnalysis* from Section 5.2 to compute the result:

```
exp :: CExp = ...
runAnalysis exp :: P ((P\Sigma Integer, Integer), Store Integer)
```

6. Monadic Parameters for Abstract Abstract Machines

In Section 5 we have demonstrated how to restore a simple collecting semantics from the monadically-parametrized semantic interface, presenting a *StorePassing* monad and an analysis with domain $\mathcal{P} ((P\Sigma Integer, Integer), Store Integer)$. More complex analyses differ from this simple analysis in a number of aspects. In this section, we will discuss how our *StorePassing* monad and *CPSInterface* and *Collecting* instances can be abstracted further to accommodate this. Specifically, we show how to control polyvariance, store representation, abstract counting, abstract garbage collection and store cloning.

⁵In the present implementation, we allow the time component to grow infinitely for simplicity, so in principle some implementation of the analysis may not terminate, which can be restricted by modifying the function *tick*.

6.1 Controlling polyvariance

So far, we have used Haskell’s *Integers* as abstract addresses in the collecting semantics. To allow experimentation with different addresses, we introduce the following class that provides a uniform view on addresses which may (optionally) incorporate context:

```
class (Ord a, Eq a) => Addressable a c | c -> a where
  tau0      :: c
  valloc    :: Var -> c -> a
  advance   :: Val a -> P\Sigma a -> c -> c
```

The type class *Addressable* has two type parameters: *a* for actual addresses and *c* for the type of the context. Contexts unambiguously define the nature of addresses. The type class contains three functions. τ_0 generates an initial context, used to instantiate the *HasInitial* class from Section 5.3.3. The function *valloc*, given a variable name and a context, allocates a new address. Finally, the function *advance* has the opportunity to internalize components of the partial state within the monad, based on the function called, the current state processed, and a given context.

A meaningful instance of the *Addressable* type class uses a combination of time-stamps, represented by lists of calls of length bounded by some *k*, and addresses that simply pair the time-stamps with variable names. This defines an analysis’ polyvariance:

```
data KTime = KCalls [CExp] deriving (Eq, Ord)
data KAddr = KBind Var KTime deriving (Eq, Ord)
```

Here is the instance of *Addressable* for this pair, relying on the auxiliary class *KCFA* for generic controlling of polyvariance:

```
class KCFA a where
  getK :: a -> Int

instance Addressable KAddr KTime where
  tau0      = KCalls []
  valloc v t = KBind v t
  advance proc (call, rho) t@(KCalls calls)
    = KCalls $ (getK t) (call : calls)
```

The *CPSInterface* instance for *StorePassing* requires a small change in order to accommodate polyvariance:

```
instance (Addressable a t) =>
  CPSInterface (StorePassing (Store a) t) a
  where
  alloc v      = gets (valloc v)
  tick proc ps = modify $ \t -> advance proc ps t
  ...         -- the rest is unchanged
```

According to the *A Posteriori* soundness theorem of Might and Manolios [17], any allocation policy for a non-deterministic abstract interpreter (which our analysis is a particular case of) leads to a *sound* abstraction of a concrete store-based collecting semantics that uses unique addresses for each allocation (e.g., integers, as in the example in Section 5.3.2). Thus, abstracting over addresses yields a family of sound abstract interpreters and requires no change in the semantics interface (Figure 2).

6.2 Abstracting over the store component

The store component is essential for efficient implementation of the analysis but can also itself be a source of valuable measurements (e.g., computing the *flows-to* information), so our next step is to make an analysis store-generic. We do so by defining a *StoreLike* class enabling creation of initial store, binding, update, and lookup as well as providing a mechanism to clean the store up:

```
class (Eq a, Lattice s, Lattice d) =>
  StoreLike a s d | s -> a, s -> d where
```

```
sigma0      :: s
bind        :: s -> a -> d -> s
replace     :: s -> a -> d -> s
fetch      :: s -> a -> d
filterStore :: s -> (a -> Bool) -> s
```

The *StoreLike* class binds together three components: addresses (*a*), the store implementation itself (*s*) and the store co-domain (*d*). To save space, we refer the reader to our public implementation for implementations of *StoreLike* instances for *Store a*.

Again, we need small changes in the *CPSInterface* instance for *StorePassing* in order to abstract over stores:

```
instance (Addressable a t, StoreLike a s (P (Val a))) =>
  CPSInterface (StorePassing s t) a
  where
  fun rho (Ref v) = lift $ getsNDSet $ flip fetch (rho ! v)
  arg rho (Ref v) = lift $ getsNDSet $ flip fetch (rho ! v)
  a -> d = lift $ modify $ \sigma -> bind sigma a (singleton d)
  ... -- the rest is unchanged
```

6.3 Controlling abstract counting

Abstract counting is a technique to track how many times an abstract resource has been allocated [18]. It provides a simple but powerful way of bounding cardinalities over abstractions. Specifically, it bounds them in a way that it enables *must*-alias analysis in imperative languages or environment analysis in higher-order languages. Counting enables additional shape analyses or predicate abstractions to be layered on top of an existing analysis [16]. This, in turn, enables an analysis’ client to perform environmental analysis to perform advanced optimization, such as super β -inlining [14].

To introduce an abstract counter for CPS, we need to make a small addition to the state-space:

$$\begin{aligned} \xi &\in \widehat{\Sigma} = \widehat{Call} \times \widehat{Env} \times \widehat{Store} \times \widehat{Count} \times \widehat{Time} \\ \hat{\mu} &\in \widehat{Count} = \widehat{Addr} \rightarrow \widehat{\mathbb{N}} \\ \hat{n} &\in \widehat{\mathbb{N}} = \{0, 1, \infty\}. \end{aligned}$$

The abstract transition relation changes correspondingly to take possible changes of $\hat{\mu}$ into account:

$$\begin{aligned} &\overbrace{(\llbracket f \ x_1 \dots x_n \rrbracket, \hat{\rho}, \hat{\sigma}, \hat{\mu}, \hat{t})}^{\xi} \rightsquigarrow (call, \hat{\rho}', \hat{\sigma}', \hat{\mu}', \hat{t}'), \text{ where} \\ &\quad \overbrace{(\llbracket \lambda (v_1 \dots v_n) \ call \rrbracket, \hat{\rho}') \in \widehat{A}(f, \hat{\rho}, \hat{\sigma})}^{clo} \\ \hat{d}_i &\in \widehat{A}(x_i, \hat{\rho}, \hat{\sigma}) & \hat{\rho}' &= \hat{\rho}'[v_i \mapsto \hat{a}_i] \\ \hat{t}' &= \widehat{tick}(clo, \xi) & \hat{\sigma}' &= \hat{\sigma} \sqcup [\hat{a}_i \mapsto \{\hat{d}_i\}] \\ \hat{a}_i &= \widehat{alloc}(v_i, \hat{t}') & \hat{\mu}' &= \hat{\mu} \oplus [\hat{a}_i \mapsto 1] \end{aligned}$$

The operator \oplus is the point-wise lifted natural abstraction of addition over $\widehat{\mathbb{N}}$. In fact, the structure of the set $\widehat{\mathbb{N}}$ can vary depending on the requirements for the analysis’ results. The only requirement is that $\widehat{\mathbb{N}}$ should be a lattice. For instance, in a degenerate case, one can turn abstract counting off by setting $\widehat{\mathbb{N}} = \{\infty\}$.

Since abstract counts are modified in tandem with the abstract store, counting and dependent enhancements like strong update can be hidden entirely within and tuned from the monad. We need no refactoring of the state-space or the semantics needed to introduce abstract counting, thanks to the introduction of the class *StoreLike* in Section 6.2. All we need to do is supply a different instance of store, namely, a “counting” one. First, we define a datatype for $\widehat{\mathbb{N}}$, make it a lattice and define abstract addition:

data $AbsNat = AZero \mid AOne \mid AMany$
deriving ($Ord, Eq, Show$)

$(\oplus) :: AbsNat \rightarrow AbsNat \rightarrow AbsNat$
 $AZero \oplus n = n$
 $n \oplus AZero = n$
 $n \oplus m = AMany$

instance $Lattice AbsNat$ **where** \dots

Second, we define a type class for abstract counter, a counting store and make the latter an instance of the former:

class $StoreLike a s d \Rightarrow ACounter a s d$ **where**
 $count :: s \rightarrow a \rightarrow AbsNat$
type $CountingStore a d = a \rightarrow (d, AbsNat)$
instance ($Ord a, Lattice d$) \Rightarrow
 $ACounter a (CountingStore a d) d$ **where**
 $count \sigma a = snd \$ \sigma ! a$

Because the counter is parameterized over addresses, it too is independent of specific semantics, and in fact can be used with any other semantics.

For a full implementation of $StoreLike a (CountingStore a d) d$, we invite the reader to puzzle through it or reference our public implementation. Once the instance is provided, a $CountingStore$ can be directly plugged into the $StorePassing$, so the abstract counting does not require any changes in the analysis logic, implicitly adding the component \widehat{Count} to $\widehat{\Sigma}$.

6.4 Controlling abstract garbage collection

Abstract garbage collection [18] is a store-sensitive analysis technique that prunes unreachable structure (exactly as ordinary garbage collection does). The net effect is an often dramatic increase in precision as well as a corresponding drop in analysis time. The technique is defined in terms of touchability of a value by a binding, the adjacency of bindings and reachability of bindings from some entity. In essence, abstract *garbage collection* stands for finding the set of reachable bindings for a particular state and restricting the domain of the store σ to solely these bindings.

The abstract bindings $\widehat{T}(x, \hat{\rho})$, touched by some abstract closure pair $(x, \hat{\rho})$ are defined as follows:

$$\begin{aligned} \widehat{T}(x, \hat{\rho}) &= \{\hat{\rho}(v) : v \in free(x)\} \\ \widehat{T}\{(x_1, \hat{\rho}_1), \dots, (x_n, \hat{\rho}_n)\} &= \widehat{T}(x_1, \hat{\rho}_1) \cup \dots \cup \widehat{T}(x_n, \hat{\rho}_n) \end{aligned}$$

We extend the notion of touching to call sites and abstract states:

$$\begin{aligned} \widehat{T}(f x_1 \dots x_n, \hat{\rho}) &= \widehat{T}(f, \hat{\rho}) \cup \widehat{T}(x_1, \hat{\rho}) \cup \dots \cup \widehat{T}(x_n, \hat{\rho}) \\ \widehat{T}(call, \hat{\rho}, \hat{\sigma}, \hat{t}) &= \widehat{T}(call, \hat{\rho}) \end{aligned}$$

and define the abstract adjacency relation for bindings:

$$\hat{a} \sim_{\hat{\sigma}} \hat{a}' \iff \hat{a}' \in \widehat{T}(\hat{\sigma}(\hat{a}))$$

Next, the abstract reachable-bindings function

$$\widehat{\mathcal{R}} : \widehat{\Sigma} \rightarrow \mathcal{P}(\widehat{Addr})$$

for a given abstract state $\hat{\zeta}$ computes the set of *reachable* bindings as all bindings we can reach from $\hat{\zeta}$ with chains of $\sim_{\hat{\sigma}}$ links:

$$\widehat{\mathcal{R}}(\hat{\zeta}) = \{\hat{a}' : \hat{a} \in \widehat{T}(\hat{\zeta}) \text{ and } \hat{a} \sim_{\hat{\sigma}}^* \hat{a}'\}.$$

We define the abstract Garbage Collection function, $\widehat{\Gamma} : \widehat{\Sigma} \rightarrow \widehat{\Sigma}$ that removes unreachable bindings from the domain of $\hat{\sigma}$:

$$\widehat{\Gamma}(call, \hat{\rho}, \hat{\sigma}, \hat{t}) = (call, \hat{\rho}, \hat{\sigma} | \widehat{\mathcal{R}}(\hat{\zeta}), \hat{t}),$$

where the vertical bar $|$ operator denotes map restriction, *i.e.*, $f | X$ is the function f defined at most over elements in the set X .

Finally, using the function $\widehat{\Gamma}$, we can define the alternate, GC abstract transition rule, $\sim_{\widehat{\Gamma}}$, so the abstract transition becomes:

$$\frac{\widehat{\Gamma}(\hat{\zeta}) \sim \hat{\zeta}'}{\hat{\zeta} \sim_{\widehat{\Gamma}} \hat{\zeta}'} \quad (\text{STEP-GC})$$

From the implementor's point of view, an abstract garbage collector modifies the *internal* part of the state, *i.e.*, an abstract store, by removing unreachable addresses. Therefore, it is natural to define abstract garbage collection abstractly as an operation in the analysis monad:

class $Monad m \Rightarrow GarbageCollector m a$ **where**
 $gc :: a \rightarrow m ()$
 $gc _ = return ()$ -- default implementation

The function gc takes a partial state as a parameter and returns a monad operation. We also supply a *default* implementation of the function gc as a no-op.

The structure of the class $StoreLike$ makes it easy to implement a garbage collector thanks to the function $filterStore$ (Section 6.2). Weaving the GC into the semantics requires only little change in the fixed point computation for $StorePassing$:

$applyStep \ step =$
 $joinWith (\lambda((\zeta, t), s) \rightarrow$
 $\dots runStateT (\mathbf{do} \{ \zeta' \leftarrow step \zeta; gc \zeta'; return \zeta' \}) t \dots)$

6.5 Controlling store-cloning

By default, the abstracted abstract machine approach to static analysis yields a heap-(/store-)cloning analysis: every state contains a store. Ordinarily, store-cloning should be reserved for situations in which the extra precision benefit the target applications.

However, for an analysis implemented this way it can take time exponential in the size of the input program when computing the reachable states of the abstracted machine [19]. The standard technique to reduce the complexity is to employ widening in the form of Shivers' single-threaded store [22]. To use a single-threaded store, we have to reconsider the abstract evaluation function itself. Instead of seeing it as a function that returns the set of reachable states, it is a function that returns a pair, consisting of a set of partial states and a single globally approximating store.

Although this change requires a significant reworking in the definition of the semantics, it is quite easy to implement in our framework, since the store is *not* a component of the program states, but rather an element supplied by an analysis monad. One can also notice that the new semantics can be captured by establishing the following Galois connection:

$$\langle \mathcal{P}(\widehat{\Sigma}_t \times \widehat{Store}), \sqsubseteq \rangle \stackrel{\gamma}{\longleftarrow \alpha} \langle \mathcal{P}(\widehat{\Sigma}_t) \times \widehat{Store}, \sqsubseteq \rangle, \quad (3)$$

where $\widehat{\Sigma}_t = \widehat{Call} \times \widehat{Env} \times \widehat{Time}$. Bounding the space of addresses (Section 6.1) implies finiteness of both lattices involved in the Galois connection (3), which means that both α and γ are computable. It is straightforward to express them in Haskell:

$alpha :: (Lattice s, Ord a, Ord g) \Rightarrow$
 $\mathcal{P}((P\Sigma a, g), s) \rightarrow \mathcal{P}(P\Sigma a, g), s)$
 $alpha = joinWith (\lambda((p, g), \sigma) \rightarrow (singleton (p, g), \sigma))$
 $gamma :: (Ord a, Ord g, Ord s) \Rightarrow$
 $\mathcal{P}(P\Sigma a, g), s) \rightarrow \mathcal{P}((P\Sigma a, g), s)$
 $gamma (states, \sigma) = Set.map (\lambda(p, g) \rightarrow ((p, g), \sigma)) states$

In words, $alpha$ combines together all per-state store components (given the stores form a lattice, hence $Lattice s$). Conversely, $gamma$ spreads the store σ among all provided states.

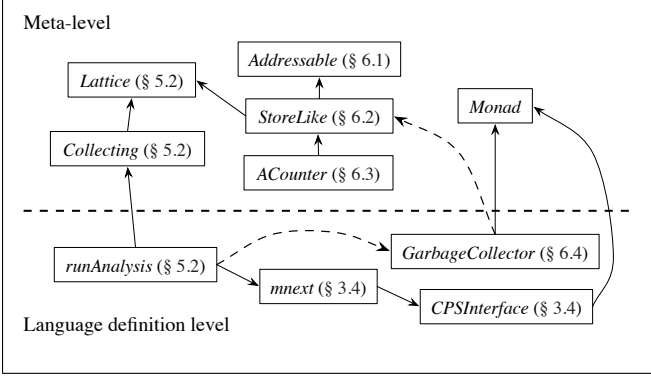


Figure 3: Overview of the framework concepts

The easiest way to construct a single-store analysis is to redefine the instance of *Collecting*, taking $F_c = \text{applyStep}' \text{ step}$:⁶

```
instance (Ord g, Ord a, Ord s, Lattice s, HasInitial g) =>
  Collecting (StorePassing s g)
  (PΣ a)
  (P (PΣ a, g), s) where
  inject a = (singleton (a, initial), ⊥)
  applyStep step = alpha ∘ (applyStep' step) ∘ gamma
```

7. Pulling It All Together

A high-level overview of the described framework is given in Figure 3. The *concepts* in the upper half are defined on the *meta-level*, i.e., can be implemented in a way so they might be reused by different languages and analyses. The lower half describes language-dependent concepts, which rely, in particular, on the fixed syntax, semantics and the structure of values. Solid arrows denote type dependencies between *instances* of the components, and dashed arrows are for the optional logical dependencies, which are, however, not strictly enforced by types. Curiously, the diagram confirms our point that store logic and abstract counting strategy are orthogonal to the analysis implementation (Sections 6.2 and 6.3).

As opposite to the abstract *concepts*, it is much harder to provide a straightforward hierarchy of actual concept *implementations* with strict “level” distinctions. For instance, one can encode the store and monad logic directly into the semantic interface definition, moving it, therefore to the bottom part of the diagram. Alternatively, one can implement the store and the monad independently from the encoding of the semantic interface, just as we did in Section 3, getting as close as possible to the conceptual decomposition. This approach pays off when composing the resulting analysis, as we demonstrate in the following section.

8. Further Examples: k -CFA Family

In this section, we expand our original examples from Section 3 to a family of k -CFA-based abstract interpreters, following the recipe described in Sections 5 and 6. We omit a few tedious implementation details for the sake of brevity; the full development and test programs can be found in the accompanying code repository:

<http://github.com/ilyasergey/monadic-cfa>

⁶We use *applyStep'* to refer to the definition of *applyStep* for a domain with per-state store. It is slightly different from the actual implementation in Haskell, as one is required to wrap the “per-state store”-version of *StorePassing* into a separate datatype in order to make it possible for a type checker to distinguish between two versions of *applyStep*. We elaborate more on this in Section 8.

8.1 A simple abstract interpreter for k -CFA

First, we fix the k -degree of the analysis by instantiating the class *KCFA* from Section 6.1 for $k = 1$:

```
instance KCFA KTime where
  getK = const 1
```

Second, we define the analysis function *analyseKCFA* by employing the *runAnalysis* function from Section 5.2:

```
analyseKCFA :: CExp →
  P ((PΣ KAddr, KTime), Store KAddr)
analyseKCFA = runAnalysis
```

The refined type of *analyseKCFA* is of particular interest. After flattening the tuples, one can see that it reflects the abstract domain of the analysis: abstract states coupled with timestamps and abstract stores.

8.2 An abstract interpreter with a shared store

As a next step, we apply the widening strategy via a shared store, as described in Section 6.5. We directly use the definitions of *alpha* and *gamma* and redefine the instance of *Collecting* by providing a new implementation of the *applyStep* function (Section 5.2). In order to overcome Haskell’s conventions for type resolution in the case of the function *applyStep*, we might need to define a *wrapper* record type for the analysis result.

```
newtype Wrap a g s =
  Wrap { unWrap :: P ((PΣ a, g), s) } deriving Lattice
```

The definition of *applyStep* for store-sharing widening looks as follows:

```
applyStep step =
  alpha ∘ unWrap ∘ applyStep step ∘ Wrap ∘ gamma
```

Note, the implementation is *not* recursive, but the inner call to *applyStep* operates with a different domain. The definition of the analysis function has the same implementation, thanks to the type class-based polymorphism [25], however, its return type is conceptually different, as it accounts for a set of states, coupled with a *single* store:

```
analyseShared :: CExp →
  P (PΣ KAddr, KTime), Store KAddr)
analyseShared = runAnalysis
```

8.3 An abstract interpreter with a counting store

As the analysis is parametrized with the store explicitly, its instrumentation with a counting machinery is trivial: we just replace the second component of the result of the single-store-passing analysis with a specialized counting store (Section 6.3).

```
type KCFACountingStore =
  CountingStore KAddr (P (Val KAddr))
analyseWithCount :: CExp →
  P (PΣ KAddr, KTime), KCFACountingStore)
analyseWithCount = runAnalysis
```

9. Related Work and Conclusion

We have illustrated a systematic method for transforming a *concrete* semantics into a monadically-parameterized machine, such that the monad determines the classical properties of an *abstract* analysis. Our work is situated firmly in the abstract interpretation tradition established by Cousot and Cousot [4, 5].

Following a series of complex power-domain constructions, Hudak and Young [10] devised simpler set-based collecting interpretations for both first-order and higher-order functional languages. They furthermore outline how to modify their approach to express *relational* properties. This requires generalizing their collecting interpretation to sets of value-environment pairs, reminiscent of our sets-of-states starting point. Like Hudak and Young, our framework can thereby describe relational properties.

The way we factor the concrete semantics is very much in the spirit of Nielson’s two-level meta-language [20]. Nielson proposes an abstract interpretation framework based on the idea of decomposing a *denotational* language definition into two parts: a *core semantics*, containing semantic rules and their types, but leaving some function symbols and domain names uninterpreted, and an *interpretation* that fills out the missing definition of function symbol and domain names, thereby allowing alternate non-standard interpretations in addition to the “standard” semantics defining the meaning of programs.

Such a decomposition is formulated in terms of a *two-level* met-language, where some types are considered to be ‘dynamically’-interpreted,⁷ and ‘dynamic’ functions symbols are represented by combinators, closed over variables of dynamically-interpreted types. This makes it possible to define the “meaning” of dynamic types and combinators in different ways to express a lazy standard semantics, detection of signs, strictness, and liveness on top of the same semantics interface. Also, Nielson further illustrates that both forward and backward analyses, independent-attribute and relational methods can be formulated in terms of the same *core semantics* given different *interpretations*. In contrast, our work focuses on abstractions of the *operational* small-step collecting semantics, since it delivers data for most of the interesting analyses in the setting of a higher-order language. In a sense, the functions from the class *CPSInterface*, such as *fun*, *arg* and others, play an analogous role to Nielson’s dynamic function symbols, and type parameters such as *a*, *s* or *g* are the dynamic types.

Our framework gains from giving function symbols monadic types: unlike Nielson’s framework, where all information about properties of interest is captured by the treatment of values of specific types, our implementation also allows us to track *temporal* properties of execution and tweak interpretation properties depending on the context of execution. Moreover, the sharing of stores as a widening strategy, is trivial to implement in our decomposition. The *well-formedness* of an interpretation in our framework is ensured by the type system of the host language (Haskell, in our case).

The core semantics we present using monads is of a more denotational flavour, as it is expressed by the *CPSInterface* type class. Recent work by Filinski, however, demonstrates that a complementary, operational representation, is possible using *reflection* and *reification* [8]. This correspondence could be used to translate the monad-based definitions of the semantics functions into specialized operational rules, which were usually hand-crafted [23].

Our publicly available implementation indicates the robustness of the approach, allowing re-use of multiple semantic aspects between different analyses and semantic formalisms.

Acknowledgements We wish to thank Olivier Danvy for discussions on Nielson’s work as well as for his hospitality during Sergey and Might’s visit to Aarhus University in December 2011, where the idea of the work was discussed for the first time. We are grateful to the PLDI 2013 reviewers for their excellent feedback. The work of Ilya Sergey has been partially supported by EU Marie Curie CO-FUND Action 291803 “Amarout-II Europe”. This research is partially funded by the Research Foundation - Flanders (FWO), and by

⁷*I.e.*, those, which can be given a specific interpretation as a domain or a lattice.

the Research Fund KU Leuven. Dominique Devriese holds a Ph.D. fellowship of the Research Foundation - Flanders (FWO). Different parts of Might’s effort on this work were partially supported by the DARPA programs APAC and CRASH.

References

- [1] M. S. Ager, O. Danvy, and J. Midtgaard. A functional correspondence between monadic evaluators and abstract machines for languages with computational effects. *Theor. Comput. Sci.*, 342(1):149–172, 2005.
- [2] M. Biernacka and O. Danvy. A syntactic correspondence between context-sensitive calculi and abstract machines. *Theor. Comput. Sci.*, 375(1-3):76–108, 2007.
- [3] M. M. T. Chakravarty, G. Keller, S. L. P. Jones, and S. Marlow. Associated types with class. In *POPL*, 2005.
- [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*, 1977.
- [5] P. Cousot and R. Cousot. Systematic design of program analysis frameworks. In *POPL*, 1979.
- [6] O. Danvy. A new one-pass transformation into monadic normal form. In *CC*, volume 2622 of *LNCS*, 2003.
- [7] O. Danvy. Defunctionalized interpreters for programming languages. In *ICFP*, 2008.
- [8] A. Filinski. Monads in action. In *POPL*, 2010.
- [9] C. Flanagan, A. Sabry, B. F. Duba, and M. Felleisen. The essence of compiling with continuations. In *PLDI*, 1993.
- [10] P. Hudak and J. Young. Collecting interpretations of expressions. *ACM Trans. Prog. Lang. Syst.*, 13(2):269–290, 1991.
- [11] M. P. Jones. Type Classes with Functional Dependencies. In *ESOP*, volume 1782 of *LNCS*, 2000.
- [12] A. Lakhota, D. R. Boccardo, A. Singh, and A. Manacero, Jr. Context-sensitive analysis of obfuscated x86 executables. In *PEPM*, 2010.
- [13] S. Liang, P. Hudak, and M. Jones. Monad transformers and modular interpreters. In *POPL*, 1995.
- [14] M. Might. *Environment analysis of higher-order languages*. PhD thesis, Georgia Institute of Technology, 2007.
- [15] M. Might. Abstract interpreters for free. In *SAS*, volume 6337 of *LNCS*, 2010.
- [16] M. Might. Shape analysis in the absence of pointers and structure. In *VMCAI*, volume 5944 of *LNCS*, 2010.
- [17] M. Might and P. Manolios. A posteriori soundness for non-deterministic abstract interpretations. In *VMCAI*, volume 5403 of *LNCS*, 2009.
- [18] M. Might and O. Shivers. Improving flow analyses via Γ CFA: abstract garbage collection and counting. In *ICFP*, 2006.
- [19] M. Might, Y. Smaragdakis, and D. Van Horn. Resolving and exploiting the *k*-CFA paradox: illuminating functional vs. object-oriented program analysis. In *PLDI*, 2010.
- [20] F. Nielson. Two-level semantics and abstract interpretation. *Theor. Comput. Sci.*, 69(2):117–242, 1989.
- [21] T. Schrijvers and B. C. Oliveira. Monads, zippers and views: virtualizing the monad stack. In *ICFP*, 2011.
- [22] O. G. Shivers. *Control-flow analysis of higher-order languages or taming lambda*. PhD thesis, Carnegie Mellon University, 1991.
- [23] D. Van Horn and M. Might. Abstracting abstract machines. In *ICFP*, 2010.
- [24] D. Van Horn and M. Might. Abstracting abstract machines: a systematic approach to higher-order program analysis. *Commun. ACM*, 54(9):101–109, 2011.
- [25] P. Wadler and S. Blott. How to make ad-hoc polymorphism less ad-hoc. In *POPL*, 1989.