

Abstracting Abstract Machines

David Van Horn *
Northeastern University
dvanhorn@ccs.neu.edu

Matthew Might
University of Utah
might@cs.utah.edu

Abstract

We describe a derivational approach to abstract interpretation that yields novel and transparently sound static analyses when applied to well-established abstract machines. To demonstrate the technique and support our claim, we transform the CEK machine of Felleisen and Friedman, a lazy variant of Krivine's machine, and the stack-inspecting CM machine of Clements and Felleisen into abstract interpretations of themselves. The resulting analyses bound temporal ordering of program events; predict return-flow and stack-inspection behavior; and approximate the flow and evaluation of by-need parameters. For all of these machines, we find that a series of well-known concrete machine refactorings, plus a technique we call store-allocated continuations, leads to machines that abstract into static analyses simply by bounding their stores. We demonstrate that the technique scales up uniformly to allow static analysis of realistic language features, including tail calls, conditionals, side effects, exceptions, first-class continuations, and even garbage collection.

Categories and Subject Descriptors F.3.2 [*Logics and Meanings of Programs*]: Semantics of Programming Languages—Program analysis, Operational semantics; F.4.1 [*Mathematical Logic and Formal Languages*]: Mathematical Logic—Lambda calculus and related systems

General Terms Languages, Theory

Keywords abstract machines, abstract interpretation

1. Introduction

Abstract machines such as the CEK machine and Krivine's machine are first-order state transition systems that represent the core of a real language implementation. Semantics-based program analysis, on the other hand, is concerned with safely approximating intensional properties of such a machine as it runs a program. It seems natural then to want to systematically derive analyses from machines to approximate the core of realistic run-time systems.

Our goal is to develop a technique that enables direct abstract interpretations of abstract machines by methods for transforming a given machine description into another that computes its finite approximation.

* Supported by the National Science Foundation under grant 0937060 to the Computing Research Association for the CIFellow Project.

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.

ICFP'10, September 27–29, 2010, Baltimore, Maryland, USA.
Copyright © 2010 ACM 978-1-60558-794-3/10/09. . . $10.00

We demonstrate that the technique of refactoring a machine with **store-allocated continuations** allows a direct structural abstraction[1] by bounding the machine's store. Thus, we are able to convert semantic techniques used to model language features into static analysis techniques for reasoning about the behavior of those very same features. By abstracting well-known machines, our technique delivers static analyzers that can reason about by-need evaluation, higher-order functions, tail calls, side effects, stack structure, exceptions and first-class continuations.

The basic idea behind store-allocated continuations is not new. SML/NJ has allocated continuations in the heap for well over a decade [28]. At first glance, modeling the program stack in an abstract machine with store-allocated continuations would not seem to provide any real benefit. Indeed, for the purpose of defining the meaning of a program, there is no benefit, because the meaning of the program does not depend on the stack-implementation strategy. Yet, a closer inspection finds that store-allocating continuations eliminate recursion from the definition of the state-space of the machine. With no recursive structure in the state-space, an abstract machine becomes eligible for conversion into an abstract interpreter through a simple structural abstraction.

To demonstrate the applicability of the approach, we derive abstract interpreters of:

- a call-by-value λ-calculus with state and control based on the CESK machine of Felleisen and Friedman [13],
- a call-by-need λ-calculus based on a tail-recursive, lazy variant of Krivine's machine derived by Ager, Danvy and Midtgaard [1], and
- a call-by-value λ-calculus with stack inspection based on the CM machine of Clements and Felleisen [3];

and use abstract garbage collection to improve precision [25].

Overview

In Section 2, we begin with the CEK machine and attempt a structural abstract interpretation, but find ourselves blocked by two recursive structures in the machine: environments and continuations. We make three refactorings to:

1. store-allocate bindings,
2. store-allocate continuations, and
3. time-stamp machine states;

resulting in the CESK, CESK*, and time-stamped CESK* machines, respectively. The time-stamps encode the history (context) of the machine's execution and facilitate context-sensitive abstractions. We then demonstrate that the time-stamped machine abstracts directly into a parameterized, sound and computable static analysis.

[1] A structural abstraction distributes component-, point-, and member-wise.

In Section 3, we replay this process (slightly abbreviated) with a lazy variant of Krivine's machine to arrive at a static analysis of by-need programs.

In Section 4, we incorporate conditionals, side effects, exceptions, first-class continuations, and garbage collection.

In Section 6, we abstract the CM (continuation-marks) machine to produce an abstract interpretation of stack inspection.

In Section 7, we widen the abstract interpretations with a single-threaded "global" store to accelerate convergence. For some of our analyzers, this widening results in polynomial-time algorithms and connects them back to known analyses.

2. From CEK to the abstract CESK*

In this section, we start with a traditional machine for a programming language based on the call-by-value λ-calculus, and gradually derive an abstract interpretation of this machine. The outline followed in this section covers the basic steps for systematically deriving abstract interpreters that we follow throughout the rest of the paper.

To begin, consider the following language of expressions:[2]

$$\begin{array}{ll} e \in \mathit{Exp} ::= x \mid (ee) \mid (\lambda x.e) \\ x \in \mathit{Var} & \text{a set of identifiers.} \end{array}$$

A standard machine for evaluating this language is the CEK machine of Felleisen and Friedman [12], and it is from this machine we derive the abstract semantics—a computable approximation of the machine's behavior. Most of the steps in this derivation correspond to well-known machine transformations and real-world implementation techniques—and most of these steps are concerned only with the *concrete machine*; a very simple abstraction is employed only at the very end.

The remainder of this section is outlined as follows: we present the CEK machine, to which we add a store, and use it to allocate variable bindings. This machine is just the CESK machine of Felleisen and Friedman [13]. From here, we further exploit the store to allocate continuations, which corresponds to a well-known implementation technique used in functional language compilers [28]. We then abstract *only the store* to obtain a framework for the sound, computable analysis of programs.

2.1 The CEK machine

A standard approach to evaluating programs is to rely on a Curry-Feys-style Standardization Theorem, which says roughly: if an expression e reduces to e' in, *e.g.*, the call-by-value λ-calculus, then e reduces to e' in a canonical manner. This canonical manner thus determines a state machine for evaluating programs: a standard reduction machine.

To define such a machine for our language, we define a grammar of evaluation contexts and notions of reduction (*e.g.*, β_v). An evaluation context is an expression with a "hole" in it. For left-to-right evaluation order, we define evaluation contexts E as:

$$E ::= [\,] \mid (Ee) \mid (vE).$$

[2] *Fine print on syntax:* As is often the case in program analysis where *semantic values* are approximated using *syntactic phrases* of the program under analysis, we would like to be able to distinguish different syntactic occurrences of otherwise identical expressions within a program. Informally, this means we want to track the source location of expressions. Formally, this is achieved by *labeling* expressions and assuming all labels within a program are distinct:

$$\begin{array}{ll} e \in \mathit{Exp} ::= x^\ell \mid (ee)^\ell \mid (\lambda x.e)^\ell \\ \ell \in \mathit{Lab} & \text{an infinite set of labels.} \end{array}$$

However, we judiciously omit labels whenever they are irrelevant and doing so improves the clarity of the presentation. Consequently, they appear only in Sections 2.7 and 7, which are concerned with k-CFA.

$$\varsigma \longmapsto_{CEK} \varsigma'$$

$\langle x, \rho, \kappa\rangle$	$\langle v, \rho', \kappa\rangle$ where $\rho(x) = (v, \rho')$
$\langle (e_0 e_1), \rho, \kappa\rangle$	$\langle e_0, \rho, \mathbf{ar}(e_1, \rho, \kappa)\rangle$
$\langle v, \rho, \mathbf{ar}(e, \rho', \kappa)\rangle$	$\langle e, \rho', \mathbf{fn}(v, \rho, \kappa)\rangle$
$\langle v, \rho, \mathbf{fn}((\lambda x.e), \rho', \kappa)\rangle$	$\langle e, \rho'[x \mapsto (v, \rho)], \kappa\rangle$

Figure 1. The CEK machine.

An expression is either a value or uniquely decomposable into an evaluation context and redex. The standard reduction machine is:

$$E[e] \longmapsto_{\beta_v} E[e'], \text{ if } e\ \beta_v\ e'.$$

However, this machine does not shed much light on a realistic implementation. At each step, the machine traverses the entire source of the program looking for a redex. When found, the redex is reduced and the contractum is plugged back in the hole, then the process is repeated.

Abstract machines such as the CEK machine, which are derivable from standard reduction machines, offer an extensionally equivalent but more realistic model of evaluation that is amenable to efficient implementation. The CEK is environment-based; it uses environments and closures to model substitution. It represents evaluation contexts as *continuations*, an inductive data structure that models contexts in an inside-out manner. The key idea of machines such as the CEK is that the whole program need not be traversed to find the next redex, consequently the machine integrates the process of plugging a contractum into a context and finding the next redex.

States of the CEK machine [12] consist of a control string (an expression), an environment that closes the control string, and a continuation:

$$\begin{array}{lll} \varsigma \in \Sigma & = & \mathit{Exp} \times \mathit{Env} \times \mathit{Kont} \\ v \in \mathit{Val} & ::= & (\lambda x.e) \\ \rho \in \mathit{Env} & = & \mathit{Var} \to_{\text{fin}} \mathit{Val} \times \mathit{Env} \\ \kappa \in \mathit{Kont} & ::= & \mathbf{mt} \mid \mathbf{ar}(e, \rho, \kappa) \mid \mathbf{fn}(v, \rho, \kappa). \end{array}$$

States are identified up to consistent renaming of bound variables.

Environments are finite maps from variables to closures. Environment extension is written $\rho[x \mapsto (v, \rho')]$.

Evaluation contexts E are represented (inside-out) by continuations as follows: $[\,]$ is represented by $\mathbf{mt}$; $E[(v[\,])]$ is represented by $\mathbf{ar}(e', \rho, \kappa)$ where ρ closes e' to represent e and κ represents E; $E[(v[\,])]$ is represented by $\mathbf{fn}(v', \rho, \kappa)$ where ρ closes v' to represent v and κ represents E.

The transition function for the CEK machine is defined in Figure 1 (we follow the textbook treatment of the CEK machine [11, page 102]). The initial machine state for a closed expression e is given by the inj function:

$$inj_{CEK}(e) = \langle e, \emptyset, \mathbf{mt}\rangle.$$

Typically, an evaluation function is defined as a partial function from closed expressions to answers:

$$eval'_{CEK}(e) = (v, \rho) \text{ if } inj(e) \longmapsto\!\!\!\!\to_{CEK} \langle v, \rho, \mathbf{mt}\rangle.$$

This gives an extensional view of the machine, which is useful, *e.g.*, to prove correctness with respect to a canonical evaluation function such as one defined by standard reduction or compositional valuation. However for the purposes of program analysis, we are concerned more with the intensional aspects of the machine. As such, we define the meaning of a program as the (possibly infinite) set of reachable machine states:

$$eval_{CEK}(e) = \{\varsigma \mid inj(e) \longmapsto\!\!\!\!\to_{CEK} \varsigma\}.$$

Deciding membership in the set of reachable machine states is not possible due to the halting problem. The goal of abstract interpretation, then, is to construct a function, $aval_{\widehat{CEK}}$, that is a sound and computable approximation to the $eval_{CEK}$ function.

We can do this by constructing a machine that is similar in structure to the CEK machine: it is defined by an *abstract state transition* relation $(\longmapsto_{\widehat{CEK}}) \subseteq \hat{\Sigma} \times \hat{\Sigma}$, which operates over *abstract states*, $\hat{\Sigma}$, which approximate the states of the CEK machine, and an abstraction map $\alpha : \Sigma \to \hat{\Sigma}$ that maps concrete machine states into abstract machine states.

The abstract evaluation function is then defined as:

$$aval_{\widehat{CEK}}(e) = \{\hat{\varsigma} \mid \alpha(inj(e)) \longmapsto\!\!\!\!\!\rightarrow_{\widehat{CEK}} \hat{\varsigma}\}.$$

1. We achieve *decidability* by constructing the approximation in such a way that the state-space of the abstracted machine is finite, which guarantees that for any closed expression e, the set $aval(e)$ is finite.
2. We achieve *soundness* by demonstrating the abstracted machine transitions preserve the abstraction map, so that if $\varsigma \longmapsto \varsigma'$ and $\alpha(\varsigma) \sqsubseteq \hat{\varsigma}$, then there exists an abstract state $\hat{\varsigma}'$ such that $\hat{\varsigma} \longmapsto \hat{\varsigma}'$ and $\alpha(\varsigma') \sqsubseteq \hat{\varsigma}'$.

A first attempt at abstract interpretation: A simple approach to abstracting the machine's state space is to apply a *structural abstract interpretation*, which lifts abstraction point-wise, element-wise, component-wise and member-wise across the structure of a machine state (*i.e.*, expressions, environments, and continuations).

The problem with the structural abstraction approach for the CEK machine is that both environments and continuations are recursive structures. As a result, the map α yields objects in an abstract state-space with recursive structure, implying the space is infinite. It is possible to perform abstract interpretation over an infinite state-space, but it requires a widening operator. A widening operator accelerates the ascent up the lattice of approximation and must guarantee convergence. It is difficult to imagine a widening operator, other than the one that jumps immediately to the top of the lattice, for these semantics.

Focusing on recursive structure as the source of the problem, a reasonable course of action is to add a level of indirection to the recursion—to force recursive structure to pass through explicitly allocated addresses. In doing so, we will unhinge recursion in a program's data structures and its control-flow from recursive structure in the state-space.

We turn our attention next to the CESK machine [10, 13], since the CESK machine eliminates recursion from one of the structures in the CEK machine: environments. In the subsequent section (Section 2.3), we will develop a CESK machine with a pointer refinement (CESK*) that eliminates the other source of recursive structure: continuations. At that point, the machine structurally abstracts via a single point of approximation: the store.

2.2 The CESK machine

The states of the CESK machine extend those of the CEK machine to include a *store*, which provides a level of indirection for variable bindings to pass through. The store is a finite map from *addresses* to *storable values* and environments are changed to map variables to addresses. When a variable's value is looked-up by the machine, it is now accomplished by using the environment to look up the variable's address, which is then used to look up the value. To bind a variable to a value, a fresh location in the store is allocated and mapped to the value; the environment is extended to map the variable to that address.

$$\varsigma \longmapsto_{CESK} \varsigma'$$

$\langle x, \rho, \sigma, \kappa\rangle$	$\langle v, \rho', \sigma, \kappa\rangle$ where $\sigma(\rho(x)) = (v, \rho')$
$\langle (e_0 e_1), \rho, \sigma, \kappa\rangle$	$\langle e_0, \rho, \sigma, \mathbf{ar}(e_1, \rho, \kappa)\rangle$
$\langle v, \rho, \sigma, \mathbf{ar}(e, \rho', \kappa)\rangle$	$\langle e, \rho', \sigma, \mathbf{fn}(v, \rho, \kappa)\rangle$
$\langle v, \rho, \sigma, \mathbf{fn}((\lambda x.e), \rho', \kappa)\rangle$	$\langle e, \rho'[x \mapsto a], \sigma[a \mapsto (v, \rho)], \kappa\rangle$ where $a \notin dom(\sigma)$

Figure 2. The CESK machine.

The state space for the CESK machine is defined as follows:

$$\begin{aligned} \varsigma \in \Sigma &= Exp \times Env \times Store \times Kont \\ \rho \in Env &= Var \to_{\text{fin}} Addr \\ \sigma \in Store &= Addr \to_{\text{fin}} Storable \\ s \in Storable &= Val \times Env \\ a, b, c \in Addr &\quad \text{an infinite set.} \end{aligned}$$

States are identified up to consistent renaming of bound variables and addresses. The transition function for the CESK machine is defined in Figure 2 (we follow the textbook treatment of the CESK machine [11, page 166]).

The initial state for a closed expression is given by the inj function, which combines the expression with the empty environment, store, and continuation:

$$inj_{CESK}(e) = \langle e, \emptyset, \emptyset, \mathbf{mt}\rangle.$$

The $eval_{CESK}$ evaluation function is defined following the template of the CEK evaluation given in Section 2.1:

$$eval_{CESK}(e) = \{\varsigma \mid inj(e) \longmapsto\!\!\!\!\!\rightarrow_{CESK} \varsigma\}.$$

Observe that for any closed expression, the CEK and CESK machines operate in lock-step: each machine transitions, by the corresponding rule, if and only if the other machine transitions.

Lemma 1 (Felleisen, [10]). $eval_{CESK}(e) \simeq eval_{CEK}(e)$.

A second attempt at abstract interpretation: With the CESK machine, half the problem with the attempted naïve abstract interpretation is solved: environments and closures are no longer mutually recursive. Unfortunately, continuations still have recursive structure. We could crudely abstract a continuation into a set of frames, losing all sense of order, but this would lead to a static analysis lacking faculties to reason about return-flow: every call would appear to return to every other call. A better solution is to refactor continuations as we did environments, redirecting the recursive structure through the store. In the next section, we explore a CESK machine with a pointer refinement for continuations.

2.3 The CESK* machine

To untie the recursive structure associated with continuations, we shift to store-allocated continuations. The $Kont$ component of the machine is replaced by a *pointer to* a continuation allocated in the store. We term the resulting machine the CESK* (control, environment, store, continuation pointer) machine. Notice the store now maps to denotable values *and* continuations:

$$\begin{aligned} \varsigma \in \Sigma &= Exp \times Env \times Store \times Addr \\ s \in Storable &= Val \times Env + Kont \\ \kappa \in Kont &::= \mathbf{mt} \mid \mathbf{ar}(e, \rho, a) \mid \mathbf{fn}(v, \rho, a). \end{aligned}$$

The revised machine is defined in Figure 3 and the initial machine state is defined as:

$$inj_{CESK^\star}(e) = \langle e, \emptyset, [a_0 \mapsto \mathbf{mt}], a_0\rangle.$$

$\varsigma \longmapsto_{CESK^\star} \varsigma'$, where $\kappa = \sigma(a), b \notin dom(\sigma)$

$\langle x, \rho, \sigma, a\rangle$	$\langle v, \rho', \sigma, a\rangle$ where $(v, \rho') = \sigma(\rho(x))$
$\langle (e_0 e_1), \rho, \sigma, a\rangle$	$\langle e_0, \rho, \sigma[b \mapsto \mathbf{ar}(e_1, \rho, a)], b\rangle$
$\langle v, \rho, \sigma, a\rangle$ if $\kappa = \mathbf{ar}(e, \rho', c)$	$\langle e, \rho', \sigma[b \mapsto \mathbf{fn}(v, \rho, c)], b\rangle$
$\langle v, \rho, \sigma, a\rangle$ if $\kappa = \mathbf{fn}((\lambda x.e), \rho', c)$	$\langle e, \rho'[x \mapsto b], \sigma[b \mapsto (v, \rho)], c\rangle$

Figure 3. The CESK* machine.

$\varsigma \longmapsto_{CESK^\star_t} \varsigma'$, where $\kappa = \sigma(a), b = alloc(\varsigma), u = tick(t)$

$\langle x, \rho, \sigma, a, t\rangle$	$\langle v, \rho', \sigma, a, u\rangle$ where $(v, \rho') = \sigma(\rho(x))$
$\langle (e_0 e_1), \rho, \sigma, a, t\rangle$	$\langle e_0, \rho, \sigma[b \mapsto \mathbf{ar}(e_1, \rho, a)], b, u\rangle$
$\langle v, \rho, \sigma, a, t\rangle$ if $\kappa = \mathbf{ar}(e, \rho, c)$	$\langle e, \rho, \sigma[b \mapsto \mathbf{fn}(v, \rho, c)], b, u\rangle$
$\langle v, \rho, \sigma, a, t\rangle$ if $\kappa = \mathbf{fn}((\lambda x.e), \rho', c)$	$\langle e, \rho'[x \mapsto b], \sigma[b \mapsto (v, \rho)], c, u\rangle$

Figure 4. The time-stamped CESK* machine.

The evaluation function (not shown) is defined along the same lines as those for the CEK (Section 2.1) and CESK (Section 2.2) machines. Like the CESK machine, it is easy to relate the CESK* machine to its predecessor; from corresponding initial configurations, these machines operate in lock-step:

Lemma 2. $eval_{CESK^\star}(e) \simeq eval_{CESK}(e)$.

Addresses, abstraction and allocation: The CESK* machine, as defined in Figure 3, nondeterministically chooses addresses when it allocates a location in the store, but because machines are identified up to consistent renaming of addresses, the transition system remains deterministic.

Looking ahead, an easy way to bound the state-space of this machine is to bound the set of addresses.[3] But once the store is finite, locations may need to be reused and when multiple values are to reside in the same location; the store will have to soundly approximate this by *joining* the values.

In our concrete machine, all that matters about an allocation strategy is that it picks an unused address. In the abstracted machine however, the strategy *may have to re-use previously allocated addresses*. The abstract allocation strategy is therefore crucial to the design of the analysis—it indicates when finite resources should be doled out and decides when information should deliberately be lost in the service of computing within bounded resources. In essence, the allocation strategy is the heart of an analysis (allocation strategies corresponding to well-known analyses are given in Section 2.7.)

For this reason, concrete allocation deserves a bit more attention in the machine. An old idea in program analysis is that dynamically allocated storage can be represented by the state of the computation at allocation time [18, 22, Section 1.2.2]. That is, allocation strategies can be based on a (representation) of the machine history. These representations are often called *time-stamps*.

A common choice for a time-stamp, popularized by Shivers [29], is to represent the history of the computation as *contours*, finite strings encoding the calling context. We present a concrete machine that uses general time-stamp approach and is parameterized by a choice of $tick$ and $alloc$ functions. We then instantiate $tick$ and $alloc$ to obtain an abstract machine for computing a k-CFA-style analysis using the contour approach.

2.4 The time-stamped CESK* machine

The machine states of the time-stamped CESK* machine include a *time* component, which is intentionally left unspecified:

$$t, u \in Time$$
$$\varsigma \in \Sigma = Exp \times Env \times Store \times Addr \times Time.$$

The machine is parameterized by the functions:

$$tick : \Sigma \to Time \qquad alloc : \Sigma \to Addr.$$

[3] A finite number of addresses leads to a finite number of environments, which leads to a finite number of closures and continuations, which in turn, leads to a finite number of stores, and finally, a finite number of states.

The $tick$ function returns the next time; the $alloc$ function allocates a fresh address for a binding or continuation. We require of $tick$ and $alloc$ that for all t and ς, $t \sqsubset tick(t)$ and $alloc(\varsigma) \notin \sigma$ where $\varsigma = \langle _, _, \sigma, _, _\rangle$.

The time-stamped CESK* machine is defined in Figure 4. Note that occurrences of ς on the right-hand side of this definition are implicitly bound to the state occurring on the left-hand side. The initial machine state is defined as:

$$inj_{CESK^\star_t}(e) = \langle e, \emptyset, [a_0 \mapsto \mathbf{mt}], a_0, t_0\rangle.$$

Satisfying definitions for the parameters are:

$$Time = Addr = \mathbb{Z}$$
$$a_0 = t_0 = 0 \quad tick\langle _, _, _, _, t\rangle = t + 1 \quad alloc\langle _, _, _, _, t\rangle = t.$$

Under these definitions, the time-stamped CESK* machine operates in lock-step with the CESK* machine, and therefore with the CESK and CEK machines as well.

Lemma 3. $eval_{CESK^\star_t}(e) \simeq eval_{CESK^\star}(e)$.

The time-stamped CESK* machine forms the basis of our abstracted machine in the following section.

2.5 The abstract time-stamped CESK* machine

As alluded to earlier, with the time-stamped CESK* machine, we now have a machine ready for direct abstract interpretation via a single point of approximation: the store. Our goal is a machine that resembles the time-stamped CESK* machine, but operates over a finite state-space and it is allowed to be nondeterministic. Once the state-space is finite, the transitive closure of the transition relation becomes computable, and this transitive closure constitutes a static analysis. Buried in a path through the transitive closure is a (possibly infinite) traversal that corresponds to the concrete execution of the program.

The abstracted variant of the time-stamped CESK* machine comes from bounding the address space of the store and the number of times available. By bounding these sets, the state-space becomes finite,[4] but for the purposes of soundness, an entry in the store may be forced to hold several values simultaneously:

$$\hat{\sigma} \in \widehat{Store} = Addr \to_{\text{fin}} \mathcal{P}(Storable).$$

Hence, stores now map an address to a *set* of storable values rather than a single value. These collections of values model approximation in the analysis. If a location in the store is re-used, the new value is joined with the current set of values. When a location is dereferenced, the analysis must consider any of the values in the set as a result of the dereference.

The abstract time-stamped CESK* machine is defined in Figure 5. The (non-deterministic) abstract transition relation changes little compared with the concrete machine. We only have to modify it to account for the possibility that multiple storable values (which

[4] Syntactic sets like *Exp* are infinite, but finite for any given program.

$\hat{\varsigma} \longmapsto_{\widehat{CESK}^\star_t} \hat{\varsigma}'$, where $\kappa \in \hat{\sigma}(a), b = \widehat{alloc}(\hat{\varsigma}, \kappa), u = \widehat{tick}(t, \kappa)$

$\langle x, \rho, \hat{\sigma}, a, t\rangle$	$\langle v, \rho', \hat{\sigma}, a, u\rangle$ where $(v, \rho') \in \hat{\sigma}(\rho(x))$
$\langle (e_0 e_1), \rho, \hat{\sigma}, a, t\rangle$	$\langle e_0, \rho, \hat{\sigma} \sqcup [b \mapsto \mathbf{ar}(e_1, \rho, a)], b, u\rangle$
$\langle v, \rho, \hat{\sigma}, a, t\rangle$ if $\kappa = \mathbf{ar}(e, \rho', c)$	$\langle e, \rho', \hat{\sigma} \sqcup [b \mapsto \mathbf{fn}(v, \rho, c)], b, u\rangle$
if $\kappa = \mathbf{fn}((\lambda x.e), \rho', c)$	$\langle e, \rho'[x \mapsto b], \hat{\sigma} \sqcup [b \mapsto (v, \rho)], c, u\rangle$

Figure 5. The abstract time-stamped CESK* machine.

includes continuations) may reside together in the store, which we handle by letting the machine non-deterministically choose a particular value from the set at a given store location.

The analysis is parameterized by abstract variants of the functions that parameterized the concrete version:

$$\widehat{tick} : \hat{\Sigma} \times Kont \to Time, \quad \widehat{alloc} : \hat{\Sigma} \times Kont \to Addr.$$

In the concrete, these parameters determine allocation and stack behavior. In the abstract, they are the arbiters of precision: they determine when an address gets re-allocated, how many addresses get allocated, and which values have to share addresses.

Recall that in the concrete semantics, these functions consume states—not states and continuations as they do here. This is because in the concrete, a state alone suffices since the state determines the continuation. But in the abstract, a continuation pointer within a state may denote a multitude of continuations; however the transition relation is defined with respect to the choice of a particular one. We thus pair states with continuations to encode the choice.

The *abstract* semantics computes the set of reachable states:

$$aval_{\widehat{CESK}^\star_t}(e) = \{\hat{\varsigma} \mid \langle e, \emptyset, [a_0 \mapsto \mathbf{mt}], a_0, t_0\rangle \longmapsto\!\!\!\!\to_{\widehat{CESK}^\star_t} \hat{\varsigma}\}.$$

2.6 Soundness and computability

The finiteness of the abstract state-space ensures decidability.

Theorem 1 (Decidability of the Abstract CESK* Machine)**.** *$\hat{\varsigma} \in aval_{\widehat{CESK}^\star_t}(e)$ is decidable.*

Proof. The state-space of the machine is non-recursive with finite sets at the leaves on the assumption that addresses are finite. Hence reachability is decidable since the abstract state-space is finite. □

We have endeavored to evolve the abstract machine gradually so that its fidelity in soundly simulating the original CEK machine is both intuitive and obvious. But to formally establish soundness of the abstract time-stamped CESK* machine, we use an abstraction function, defined in Figure 6, from the state-space of the concrete time-stamped machine into the abstracted state-space.

The abstraction map over times and addresses is defined so that the parameters $\widehat{alloc}$ and $\widehat{tick}$ are sound simulations of the parameters $alloc$ and $tick$, respectively. We also define the partial order ($\sqsubseteq$) on the abstract state-space as the natural point-wise, element-wise, component-wise and member-wise lifting, wherein the partial orders on the sets Exp and $Addr$ are flat. Then, we can prove that abstract machine's transition relation simulates the concrete machine's transition relation.

Theorem 2 (Soundness of the Abstract CESK* Machine)**.** *If $\varsigma \longmapsto_{CEK} \varsigma'$ and $\alpha(\varsigma) \sqsubseteq \hat{\varsigma}$, then there exists an abstract state $\hat{\varsigma}'$, such that $\hat{\varsigma} \longmapsto_{\widehat{CESK}^\star_t} \hat{\varsigma}'$ and $\alpha(\varsigma') \sqsubseteq \hat{\varsigma}'$.*

Proof. By Lemmas 1, 2, and 3, it suffices to prove soundness with respect to $\longmapsto_{CESK^\star_t}$. Assume $\varsigma \longmapsto_{CESK^\star_t} \varsigma'$ and $\alpha(\varsigma) \sqsubseteq \hat{\varsigma}$.

$$\begin{aligned}
\alpha(e, \rho, \sigma, a, t) &= (e, \alpha(\rho), \alpha(\sigma), \alpha(a), \alpha(t)) && \text{[states]}\\
\alpha(\rho) &= \lambda x.\alpha(\rho(x)) && \text{[environments]}\\
\alpha(\sigma) &= \lambda \hat{a}. \bigsqcup_{\alpha(a)=\hat{a}} \{\alpha(\sigma(a))\} && \text{[stores]}\\
\alpha((\lambda x.e), \rho) &= ((\lambda x.e), \alpha(\rho)) && \text{[closures]}\\
\alpha(\mathbf{mt}) &= \mathbf{mt} && \text{[continuations]}\\
\alpha(\mathbf{ar}(e, \rho, a)) &= \mathbf{ar}(e, \alpha(\rho), \alpha(a))\\
\alpha(\mathbf{fn}(v, \rho, a)) &= \mathbf{fn}(v, \alpha(\rho), \alpha(a)),
\end{aligned}$$

Figure 6. The abstraction map, $\alpha : \Sigma_{CESK^\star_t} \to \hat{\Sigma}_{\widehat{CESK}^\star_t}$.

Because ς transitioned, exactly one of the rules from the definition of ($\longmapsto_{CESK^\star_t}$) applies. We split by cases on these rules. The rule for the second case is deterministic and follows by calculation. For the the remaining (nondeterministic) cases, we must show an abstract state exists such that the simulation is preserved. By examining the rules for these cases, we see that all three hinge on the abstract store in $\hat{\varsigma}$ soundly approximating the concrete store in ς, which follows from the assumption that $\alpha(\varsigma) \sqsubseteq \hat{\varsigma}$. □

2.7 A k-CFA-like abstract CESK* machine

In this section, we instantiate the time-stamped CESK* machine to obtain a contour-based machine; this instantiation forms the basis of a context-sensitive abstract interpreter with polyvariance like that found in k-CFA [29]. In preparation for abstraction, we instantiate the time-stamped machine using labeled call strings.

Inside times, we use *contours* ($Contour$), which are finite strings of call site labels that describe the current context:

$$\delta \in Contour ::= \epsilon \mid \ell\delta.$$

The labeled CESK machine transition relation must appropriately instantiate the parameters $tick$ and $alloc$ to augment the time-stamp on function call.

Next, we switch to abstract stores and bound the address space by truncating call string contours to length at most k (for k-CFA):

$$\delta \in \widehat{Contour}_k \text{ iff } \delta \in Contour \text{ and } |\delta| \leq k.$$

Combining these changes, we arrive at the instantiations for the concrete and abstract machines given in Figure 7, where the value $\lfloor\delta\rfloor_k$ is the leftmost k labels of contour δ.

Comparison to k-CFA: We say "k-CFA-like" rather than "k-CFA" because there are distinctions between the machine just described and k-CFA:

1. k-CFA focuses on "what flows where"; the ordering between states in the abstract transition graph produced by our machine produces "what flows where *and when*."
2. Standard presentations of k-CFA implicitly inline a global approximation of the store into the algorithm [29]; ours uses one store per state to increase precision at the cost of complexity. In terms of our framework, the lattice through which classical k-CFA ascends is $\mathcal{P}(Exp \times Env \times Addr) \times \widehat{Store}$, whereas our analysis ascends the lattice $\mathcal{P}\left(Exp \times Env \times \widehat{Store} \times Addr\right)$. We can explicitly inline the store to achieve the same complexity, as shown in Section 7.
3. On function call, k-CFA merges argument values together with previous instances of those arguments from the same context; our "minimalist" evolution of the abstract machine takes a

$$
\begin{aligned}
Time &= (Lab + \bullet) \times Contour \\
Addr &= (Lab + Var) \times Contour \\
t_0 &= (\bullet, \epsilon) \\
tick\langle x, _, _, _, t\rangle &= t \\
tick\langle (e_0 e_1)^\ell, _, _, _, (_, \delta)\rangle &= (\ell, \delta) \\
tick\langle v, _, \sigma, a, (\ell, \delta)\rangle &= \begin{cases} (\ell, \delta), & \text{if } \sigma(a) = \mathbf{ar}(_, _, _) \\ (\bullet, \ell\delta), & \text{if } \sigma(a) = \mathbf{fn}(_, _, _) \end{cases} \\
alloc(\langle (e_0^\ell e_1), _, _, _, (_, \delta)\rangle) &= (\ell, \delta) \\
alloc(\langle v, _, \sigma, a, (_, \delta)\rangle) &= (\ell, \delta) \text{ if } \sigma(a) = \mathbf{ar}(e^\ell, _, _) \\
alloc(\langle v, _, \sigma, a, (_, \delta)\rangle) &= (x, \delta) \text{ if } \sigma(a) = \mathbf{fn}((\lambda x.e), _, _) \\
\widehat{tick}(\langle x, _, _, _, t\rangle, \kappa) &= t \\
\widehat{tick}(\langle (e_0 e_1)^\ell, _, _, _, (_, \delta)\rangle, \kappa) &= (\ell, \delta) \\
\widehat{tick}(\langle v, _, \hat\sigma, a, (\ell, \delta)\rangle, \kappa) &= \begin{cases} (\ell, \delta), & \text{if } \kappa = \mathbf{ar}(_, _, _) \\ (\bullet, \lfloor \ell\delta \rfloor_k), & \text{if } \kappa = \mathbf{fn}(_, _, _) \end{cases} \\
\widehat{alloc}(\langle (e_0^\ell e_1), _, _, _, (_, \delta)\rangle, \kappa) &= (\ell, \delta) \\
\widehat{alloc}(\langle v, _, \hat\sigma, a, (_, \delta)\rangle, \kappa) &= (\ell, \delta) \text{ if } \kappa = \mathbf{ar}(e^\ell, _, _) \\
\widehat{alloc}(\langle v, _, \hat\sigma, a, (_, \delta)\rangle, \kappa) &= (x, \delta) \text{ if } \kappa = \mathbf{fn}((\lambda x.e), _, _)
\end{aligned}
$$

Figure 7. Instantiation for k-CFA machine.

higher-precision approach: it forks the machine for each argument value, rather than merging them immediately.

4. k-CFA does not recover explicit information about stack structure; our machine contains an explicit model of the stack for every machine state.

3. Analyzing by-need with Krivine's machine

Even though the abstract machines of the prior section have advantages over traditional CFAs, the approach we took (store-allocated continuations) yields more novel results when applied in a different context: a lazy variant of Krivine's machine. That is, we can construct an abstract interpreter that both analyzes and exploits laziness. Specifically, we present an abstract analog to a lazy and properly tail-recursive variant of Krivine's machine [19, 20] derived by Ager, Danvy, and Midtgaard [1]. The derivation from Ager *et al.*'s machine to the abstract interpreter follows the same outline as that of Section 2: we apply a pointer refinement by store-allocating continuations and carry out approximation by bounding the store.

The by-need variant of Krivine's machine considered here uses the common implementation technique of store-allocating thunks and forced values. When an application is evaluated, a thunk is created that will compute the value of the argument when forced. When a variable occurrence is evaluated, if it is bound to a thunk, the thunk is forced (evaluated) and the store is updated to the result. Otherwise if a variable occurrence is evaluated and bound to a forced value, that value is returned.

Storable values include delayed computations (thunks) $\mathbf{d}(e, \rho)$, and computed values $\mathbf{c}(v, \rho)$, which are just tagged closures. There are two continuation constructors: $\mathbf{c_1}(a, \kappa)$ is induced by a variable occurrence whose binding has not yet been forced to a value. The address a is where we want to write the given value when this continuation is invoked. The other: $\mathbf{c_2}(a, \kappa)$ is induced by an application expression, which forces the operator expression to a value. The address a is the address of the argument.

$\varsigma \longmapsto_{LK} \varsigma'$

$\langle x, \rho, \sigma, \kappa\rangle$ if $\sigma(\rho(x)) = \mathbf{d}(e, \rho')$	$\langle e, \rho', \sigma, \mathbf{c_1}(\rho(x), \kappa)\rangle$
if $\sigma(\rho(x)) = \mathbf{c}(v, \rho')$	$\langle v, \rho', \sigma, \kappa\rangle$
$\langle (e_0 e_1), \rho, \sigma, \kappa\rangle$	$\langle e_0, \rho, \sigma[a \mapsto \mathbf{d}(e_1, \rho)], \mathbf{c_2}(a, \kappa)\rangle$ where $a \notin dom(\sigma)$
$\langle v, \rho, \sigma, \mathbf{c_1}(a, \kappa)\rangle$	$\langle v, \rho, \sigma[a \mapsto \mathbf{c}(v, \rho)], \kappa\rangle$
$\langle (\lambda x.e), \rho, \sigma, \mathbf{c_2}(a, \kappa)\rangle$	$\langle e, \rho[x \mapsto a], \sigma, \kappa\rangle$

Figure 8. The LK machine.

$\hat\varsigma \longmapsto_{\widehat{LK^\star_t}} \hat\varsigma'$, where $\kappa \in \hat\sigma(a), b = \widehat{alloc}(\hat\varsigma, \kappa), u = \widehat{tick}(\hat\varsigma, \kappa)$

$\langle x, \rho, \hat\sigma, a, t\rangle$ if $\hat\sigma(\rho(x)) \ni \mathbf{d}(e, \rho')$	$\langle e, \rho', \hat\sigma \sqcup [b \mapsto \mathbf{c_1}(\rho(x), a)], b, u\rangle$
$\langle x, \rho, \hat\sigma, a, t\rangle$ if $\hat\sigma(\rho(x)) \ni \mathbf{c}(v, \rho')$	$\langle v, \rho', \hat\sigma, a, u\rangle$
$\langle (e_0 e_1), \rho, \hat\sigma, a, t\rangle$	$\langle e_0, \rho, \hat\sigma', b, u\rangle$ where $c = \widehat{alloc}(\hat\varsigma, \kappa)$, $\hat\sigma' = \hat\sigma \sqcup [c \mapsto \mathbf{d}(e_1, \rho), b \mapsto \mathbf{c_2}(c, a)]$
$\langle v, \rho, \hat\sigma, a, t\rangle$ if $\kappa = \mathbf{c_1}(a', c)$	$\langle v, \rho', \hat\sigma \sqcup [a' \mapsto \mathbf{c}(v, \rho)], c, u\rangle$
$\langle (\lambda x.e), \rho, \hat\sigma, a, t\rangle$ if $\kappa = \mathbf{c_2}(a', c)$	$\langle e, \rho'[x \mapsto a'], \hat\sigma, c, u\rangle$

Figure 9. The abstract LK* machine.

The concrete state-space is defined as follows and the transition relation is defined in Figure 8:

$$
\begin{aligned}
\varsigma \in \Sigma &= Exp \times Env \times Store \times Kont \\
s \in Storable &::= \mathbf{d}(e, \rho) \mid \mathbf{c}(v, \rho) \\
\kappa \in Kont &::= \mathbf{mt} \mid \mathbf{c_1}(a, \kappa) \mid \mathbf{c_2}(a, \kappa)
\end{aligned}
$$

When the control component is a variable, the machine looks up its stored value, which is either computed or delayed. If delayed, a $\mathbf{c_1}$ continuation is pushed and the frozen expression is put in control. If computed, the value is simply returned. When a value is returned to a $\mathbf{c_1}$ continuation, the store is updated to reflect the computed value. When a value is returned to a $\mathbf{c_2}$ continuation, its body is put in control and the formal parameter is bound to the address of the argument.

We now refactor the machine to use store-allocated continuations; storable values are extended to include continuations:

$$
\begin{aligned}
\varsigma \in \Sigma &= Exp \times Env \times Store \times Addr \\
s \in Storable &::= \mathbf{d}(e, \rho) \mid \mathbf{c}(v, \rho) \mid \kappa \\
\kappa \in Kont &::= \mathbf{mt} \mid \mathbf{c_1}(a, a) \mid \mathbf{c_2}(a, a).
\end{aligned}
$$

It is straightforward to perform a pointer-refinement of the LK machine to store-allocate continuations as done for the CESK machine in Section 2.3 and observe the lazy variant of Krivine's machine and its pointer-refined counterpart (not shown) operate in lock-step:

Lemma 4. $eval_{LK}(e) \simeq eval_{LK^\star}(e)$.

After threading time-stamps through the machine as done in Section 2.4 and defining $\widehat{tick}$ and $\widehat{alloc}$ analogously to the defi-

nitions given in Section 2.5, the pointer-refined machine abstracts directly to yield the abstract LK* machine in Figure 9.

The abstraction map for this machine is a straightforward structural abstraction similar to that given in Section 2.6 (and hence omitted). The abstracted machine is sound with respect to the LK* machine, and therefore the original LK machine.

Theorem 3 (Soundness of the Abstract LK* Machine)**.** *If $\varsigma \longmapsto_{LK} \varsigma'$ and $\alpha(\varsigma) \sqsubseteq \hat{\varsigma}$, then there exists an abstract state $\hat{\varsigma}'$, such that $\hat{\varsigma} \longmapsto_{\widehat{LK^\star_t}} \hat{\varsigma}'$ and $\alpha(\varsigma') \sqsubseteq \hat{\varsigma}'$.*

Optimizing the machine through specialization: Ager *et al.* optimize the LK machine by specializing application transitions. When the operand of an application is a variable, no delayed computation needs to be constructed, thus "avoiding the construction of space-leaky chains of thunks." Likewise, when the operand is a λ-abstraction, "we can store the corresponding closure as a computed value rather than as a delayed computation." Both of these optimizations, which conserve valuable abstract resources, can be added with no trouble, as shown in Figure 10.

$\hat{\varsigma} \longmapsto_{\widehat{LK^\star}} \hat{\varsigma}'$, where $\kappa \in \hat{\sigma}(a), b = \widehat{alloc}(\hat{\varsigma}, \kappa), u = \widehat{tick}(t)$

$\langle (ex), \rho, \hat{\sigma}, a, t\rangle$	$\langle e, \rho, \hat{\sigma} \sqcup [b \mapsto \mathbf{c_2}(\rho(x), a)], b, u\rangle$
$\langle (ev), \rho, \hat{\sigma}, a, t\rangle$	$\langle e_0, \rho, \hat{\sigma} \sqcup [b \mapsto \mathbf{c}(v, \rho), c \mapsto \mathbf{c_2}(b, a)], c, u\rangle$ where $c = \widehat{alloc}(\hat{\varsigma}, \kappa)$

Figure 10. The abstract optimized LK* machine.

Varying the machine through postponed thunk creation: Ager *et al.* also vary the LK machine by postponing the construction of a delayed computation from the point at which an application is the control string to the point at which the operator has been evaluated and is being applied. The $\mathbf{c_2}$ continuation is modified to hold, rather than the address of a delayed computation, the constituents of the computation itself:

$$\kappa \in Kont ::= \mathbf{mt} \mid \mathbf{c_1}(a, a) \mid \mathbf{c_2}(e, \rho, a).$$

The transitions for applications and functions are replaced with those in Figure 11. This allocates thunks when a function is applied, rather than when the control string is an application.

As Ager *et al.* remark, each of these variants gives rise to an abstract machine. From each of these machines, we are able to systematically derive their abstractions.

4. State and control

We have shown that store-allocated continuations make abstract interpretation of the CESK machine and a lazy variant of Krivine's machine straightforward. In this section, we want to show that the tight correspondence between concrete and abstract persists after the addition of language features such as conditionals, side effects, exceptions and continuations. We tackle each feature, and present the additional machinery required to handle each one. In most cases, the path from a canonical concrete machine to pointer-refined abstraction of the machine is so simple we only show the abstracted system. In doing so, we are arguing that this abstract machine-oriented approach to abstract interpretation represents a flexible and viable framework for building abstract interpreters.

4.1 Conditionals, mutation, and control

To handle conditionals, we extend the language with a new syntactic form, `(if e e e)`, and introduce a base value `#f`, representing false. Conditional expressions induce a new continuation form: $\mathbf{if}(e'_0, e'_1, \rho, a)$, which represents the evaluation context

$\hat{\varsigma} \longmapsto_{\widehat{LK'^\star}} \hat{\varsigma}'$, where $\kappa \in \hat{\sigma}(a), b = \widehat{alloc}(\hat{\varsigma}, \kappa), u = \widehat{tick}(t)$

$\langle (e_0 e_1), \rho, \hat{\sigma}, a\rangle$	$\langle e_0, \rho, \hat{\sigma} \sqcup [b \mapsto \mathbf{c_2}(e_1, \rho, a)], b\rangle$
$\langle (\lambda x.e), \rho, \hat{\sigma}, a\rangle$ if $\kappa = \mathbf{c_2}(e', \rho', c)$	$\langle e, \rho[x \mapsto b], \hat{\sigma} \sqcup [b \mapsto \mathbf{d}(e', \rho')], c\rangle$

Figure 11. The abstract thunk postponing LK* machine.

$\hat{\varsigma} \longmapsto_{\widehat{CESK^\star_t}} \hat{\varsigma}'$, where $\kappa \in \hat{\sigma}(a), b = \widehat{alloc}(\hat{\varsigma}, \kappa), u = \widehat{tick}(t)$

$\langle (\texttt{if } e_0\ e_1\ e_2), \rho, \hat{\sigma}, a, t\rangle$	$\langle e_0, \rho, \hat{\sigma} \sqcup [b \mapsto \mathbf{if}(e_1, e_2, \rho, a)], b, u\rangle$
$\langle \texttt{\#f}, \rho, \hat{\sigma}, a, t\rangle$ if $\kappa = \mathbf{if}(e_0, e_1, \rho', c)$	$\langle e_1, \rho', \hat{\sigma}, c, u\rangle$
$\langle v, \rho, \hat{\sigma}, a, t\rangle$ if $\kappa = \mathbf{if}(e_0, e_1, \rho', c)$, and $v \neq \texttt{\#f}$	$\langle e_0, \rho', \hat{\sigma}, c, u\rangle$
$\langle (\texttt{set!}\ x\ e), \rho, \hat{\sigma}, a, t\rangle$	$\langle e, \rho, \hat{\sigma} \sqcup [b \mapsto \mathbf{set}(\rho(x), a)], b, u\rangle$
$\langle v, \rho, \hat{\sigma}, a, t\rangle$ if $\kappa = \mathbf{set}(a', c)$	$\langle v', \rho, \hat{\sigma} \sqcup [a' \mapsto v], c, u\rangle$ where $v' \in \hat{\sigma}(a')$
$\langle (\lambda x.e), \rho, \hat{\sigma}, a, t\rangle$ if $\kappa = \mathbf{fn}(\texttt{callcc}, \rho', c)$	$\langle e, \rho[x \mapsto b], \hat{\sigma} \sqcup [b \mapsto c], c, u\rangle$ where $c = \widehat{alloc}(\hat{\varsigma}, \kappa)$
$\langle c, \rho, \hat{\sigma}, a, t\rangle$ if $\kappa = \mathbf{fn}(\texttt{callcc}, \rho', a')$	$\langle a, \rho, \hat{\sigma}, c, u\rangle$
$\langle v, \rho, \hat{\sigma}, a, t\rangle$ if $\kappa = \mathbf{fn}(c, \rho', a')$	$\langle v, \rho, \hat{\sigma}, c, u\rangle$

Figure 12. The abstract extended CESK* machine.

$E[\texttt{(if [ ] } e_0\ e_1\texttt{)}]$ where ρ closes e'_0 to represent e_0, ρ closes e'_1 to represent e_1, and a is the address of the representation of E.

Side effects are fully amenable to our approach; we introduce Scheme's `set!` for mutating variables using the `(set!` x e`)` syntax. The `set!` form evaluates its subexpression e and assigns the value to the variable x. Although `set!` expressions are evaluated for effect, we follow Felleisen *et al.* and specify `set!` expressions evaluate to the value of x before it was mutated [11, page 166]. The evaluation context $E[\texttt{(set! } x\ \texttt{[ ])}]$ is represented by $\mathbf{set}(a_0, a_1)$, where a_0 is the address of x's value and a_1 is the address of the representation of E.

First-class control is introduced by adding a new base value `callcc` which reifies the continuation as a new kind of applicable value. Denoted values are extended to include representations of continuations. Since continuations are store-allocated, we choose to represent them by address. When an address is applied, it represents the application of a continuation (reified via `callcc`) to a value. The continuation at that point is discarded and the applied address is installed as the continuation.

The resulting grammar is:

$$\begin{array}{rcl} e \in Exp & ::= & \ldots \mid \texttt{(if } e\ e\ e\texttt{)} \mid \texttt{(set!} x\ e\texttt{)} \\ \kappa \in Kont & ::= & \ldots \mid \mathbf{if}(e, e, \rho, a) \mid \mathbf{set}(a, a) \\ v \in Val & ::= & \ldots \mid \texttt{\#f} \mid \texttt{callcc} \mid a. \end{array}$$

We show only the abstract transitions, which result from store-allocating continuations, time-stamping, and abstracting the concrete transitions for conditionals, mutation, and control. The first three machine transitions deal with conditionals; here we follow the Scheme tradition of considering all non-false values as true. The fourth and fifth transitions deal with mutation.

$\varsigma \longmapsto_{CESHK} \varsigma'$

$\langle v, \rho, \sigma, \mathbf{hn}(v', \rho', \kappa, \eta), \mathbf{mt}\rangle$	$\langle v, \rho, \sigma, \eta, \kappa\rangle$
$\langle (\texttt{throw}\ v), \rho, \sigma, \mathbf{hn}((\lambda x.e), \rho', \kappa', \eta), \kappa\rangle$	$\langle e, \rho'[x \mapsto a], \sigma[a \mapsto (v, \rho)], \eta, \kappa'\rangle$ where $a \notin dom(\sigma)$
$\langle (\texttt{catch}\ e\ v), \rho, \sigma, \eta, \kappa\rangle$	$\langle e, \rho, \sigma, \mathbf{hn}(v, \rho, \kappa, \eta), \mathbf{mt}\rangle$

Figure 13. The CESHK machine.

The remaining three transitions deal with first-class control. In the first of these, `callcc` is being applied to a closure value v. The value v is then "called with the current continuation", *i.e.*, v is applied to a value that represents the continuation at this point. In the second, `callcc` is being applied to a continuation (address). When this value is applied to the reified continuation, it aborts the current computation, installs itself as the current continuation, and puts the reified continuation "in the hole". Finally, in the third, a continuation is being applied; c gets thrown away, and v gets plugged into the continuation b.

In all cases, these transitions result from pointer-refinement, time-stamping, and abstraction of the usual machine transitions.

4.2 Exceptions and handlers

To analyze exceptional control flow, we extend the CESK machine with a register to hold a stack of exception handlers. This models a reduction semantics in which we have two additional kinds of evaluation contexts:

$$\begin{array}{l} E ::= [\,] \mid (\texttt{(}Ee\texttt{)}) \mid (\texttt{(}vE\texttt{)}) \mid (\texttt{catch}\ E\ v) \\ F ::= [\,] \mid (\texttt{(}Fe\texttt{)}) \mid (\texttt{(}vF\texttt{)}) \\ H ::= [\,] \mid H[F[(\texttt{catch}\ H\ v)]], \end{array}$$

and the additional, context-sensitive, notions of reduction:

$$(\texttt{catch}\ E[(\texttt{throw}\ v)]\ v') \rightarrow (\texttt{(}v'v\texttt{)}), \quad (\texttt{catch}\ v\ v') \rightarrow v.$$

H contexts represent a stack of exception handlers, while F contexts represent a "local" continuation, *i.e.*, the rest of the computation (with respect to the hole) up to an enclosing handler, if any. E contexts represent the entire rest of the computation, including handlers.

The language is extended with expressions for raising and catching exceptions. A new kind of continuation is introduced to represent a stack of handlers. In each frame of the stack, there is a procedure for handling an exception and a (handler-free) continuation:

$$\begin{array}{l} e \in Exp \quad ::= \ldots \mid (\texttt{throw}\ v) \mid (\texttt{catch}\ e\ (\lambda x.e)) \\ \eta \in Handl ::= \mathbf{mt} \mid \mathbf{hn}(v, \rho, \kappa, \eta) \end{array}$$

An η continuation represents a stack of exception handler contexts, *i.e.*, $\mathbf{hn}(v', \rho, \kappa, \eta)$ represents $H[F[(\texttt{catch}\ [\,]\ v)]]$, where η represents H, κ represents F, and ρ closes v' to represent v.

The machine includes all of the transitions of the CESK machine extended with a η component; these transitions are omitted for brevity. The additional transitions are given in Figure 13. This presentation is based on a textbook treatment of exceptions and handlers [11, page 135].[5] The initial configuration is given by:

$$inj_{CESHK}(e) = \langle e, \emptyset, \emptyset, \mathbf{mt}, \mathbf{mt}\rangle.$$

[5] To be precise, Felleisen *et al.* present the CHC machine, a substitution based machine that uses evaluation contexts in place of continuations. Deriving the CESHK machine from it is an easy exercise.

$\varsigma \longmapsto_{CESHK^\star} \varsigma'$, where $\eta = \sigma(h), \kappa = \sigma(a), b \notin dom(\sigma)$

$\langle v, \rho, \sigma, h, a\rangle$ if $\eta = \mathbf{hn}(v', \rho', a', h')$, and $\kappa = \mathbf{mt}$	$\langle v, \rho, \sigma, h', a'\rangle$
$\langle (\texttt{throw}\ v), \rho, \sigma, h, a\rangle$ if $\eta = \mathbf{hn}((\lambda x.e), \rho', a', h')$	$\langle e, \rho'[x \mapsto b], \sigma[b \mapsto (v, \rho)], h', a'\rangle$
$\langle (\texttt{catch}\ e\ v), \rho, \sigma, h, a\rangle$	$\langle e, \rho, \sigma[b \mapsto \mathbf{hn}(v, \rho, a, h)], b, a_{\mathbf{mt}}\rangle$

Figure 14. The CESHK* machine.

$\hat{\varsigma} \longmapsto_{\widehat{CESHK}^\star_t} \hat{\varsigma}'$, where $\eta \in \hat{\sigma}(h), \kappa \in \hat{\sigma}(a), b = \widehat{alloc}(\hat{\varsigma}, \eta, \kappa)$, $u = \widehat{tick}(t)$

$\langle v, \rho, \hat{\sigma}, h, a, t\rangle$ if $\eta = \mathbf{hn}(v', \rho', a', h')$, and $\kappa = \mathbf{mt}$	$\langle v, \rho, \hat{\sigma}, h', a', u\rangle$
$\langle (\texttt{throw}\ v), \rho, \hat{\sigma}, h, a, t\rangle$ if $\eta = \mathbf{hn}((\lambda x.e), \rho', a', h')$	$\langle e, \rho'[x \mapsto b], \hat{\sigma} \sqcup [b \mapsto (v, \rho)], h', a', u\rangle$
$\langle (\texttt{catch}\ e\ v), \rho, \hat{\sigma}, h, a, t\rangle$	$\langle e, \rho, \hat{\sigma} \sqcup [b \mapsto \mathbf{hn}(v, \rho, a, h)], b, a_{\mathbf{mt}}, u\rangle$

Figure 15. The abstract CESHK* machine.

In the pointer-refined machine, the grammar of handler continuations changes to the following:

$$\eta \in Handl ::= \mathbf{mt} \mid \mathbf{hn}(v, \rho, a, h),$$

where h is used to range over addresses pointing to handler continuations. The notation $a_{\mathbf{mt}}$ means a such that $\sigma(a) = \mathbf{mt}$ in concrete case and $\mathbf{mt} \in \hat{\sigma}(a)$ in the abstract, where the intended store should be clear from context. The pointer-refined machine is given in Figure 14.

After threading time-stamps through the machine as done in Section 2.4, the machine abstracts as usual to obtain the machine in Figure 15. The only unusual step in the derivation is to observe that some machine transitions rely on a choice of *two* continuations from the store; a handler and a local continuation. Analogously to Section 2.5, we extend $\widehat{tick}$ and $\widehat{alloc}$ to take two continuation arguments to encode the choice:

$$\begin{array}{l} \widehat{tick} : \hat{\Sigma} \times Handl \times Kont \rightarrow Time, \\ \widehat{alloc} : \hat{\Sigma} \times Handl \times Kont \rightarrow Addr. \end{array}$$

5. Abstract garbage collection

Garbage collection determines when a store location has become unreachable and can be re-allocated. This is significant in the abstract semantics because an address may be allocated to multiple values due to finiteness of the address space. Without garbage collection, the values allocated to this common address must be joined, introducing imprecision in the analysis (and inducing further, perhaps spurious, computation). By incorporating garbage collection in the abstract semantics, the location may be proved to be unreachable and safely *overwritten* rather than joined, in which case no imprecision is introduced.

Like the rest of the features addressed in this paper, we can incorporate abstract garbage collection into our static analyzers

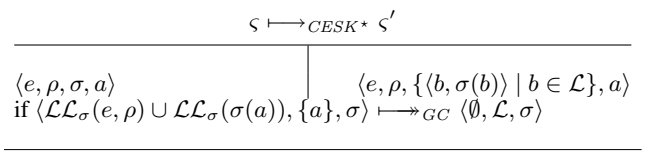


Figure 16. The GC transition for the CESK* machine.

by a straightforward pointer-refinement of textbook accounts of concrete garbage collection, followed by a finite store abstraction.

Concrete garbage collection is defined in terms of a GC machine that computes the reachable addresses in a store [11, page 172]:

$$\langle \mathcal{G}, \mathcal{B}, \sigma\rangle \longmapsto_{GC} \langle (\mathcal{G} \cup \mathcal{LL}_\sigma(\sigma(a)) \setminus (\mathcal{B} \cup \{a\})), \mathcal{B} \cup \{a\}, \sigma\rangle \text{ if } a \in \mathcal{G}.$$

This machine iterates over a set of reachable but unvisited "grey" locations $\mathcal{G}$. On each iteration, an element is removed and added to the set of reachable and visited "black" locations $\mathcal{B}$. Any newly reachable and unvisited locations, as determined by the "live locations" function $\mathcal{LL}_\sigma$, are added to the grey set. When there are no grey locations, the black set contains all reachable locations. Everything else is garbage.

The live locations function computes a set of locations which may be used in the store. Its definition will vary based on the particular machine being garbage collected, but the definition appropriate for the CESK* machine of Section 2.3 is

$$\begin{aligned}
\mathcal{LL}_\sigma(e) &= \emptyset \\
\mathcal{LL}_\sigma(e,\rho) &= \mathcal{LL}_\sigma(\rho|\mathbf{fv}(e)) \\
\mathcal{LL}_\sigma(\rho) &= rng(\rho) \\
\mathcal{LL}_\sigma(\mathbf{mt}) &= \emptyset \\
\mathcal{LL}_\sigma(\mathbf{fn}(v,\rho,a)) &= \{a\} \cup \mathcal{LL}_\sigma(v,\rho) \cup \mathcal{LL}_\sigma(\sigma(a)) \\
\mathcal{LL}_\sigma(\mathbf{ar}(e,\rho,a)) &= \{a\} \cup \mathcal{LL}_\sigma(e,\rho) \cup \mathcal{LL}_\sigma(\sigma(a)).
\end{aligned}$$

We write $\rho|\mathbf{fv}(e)$ to mean ρ restricted to the domain of free variables in e. We assume the least-fixed-point solution in the calculation of the function $\mathcal{LL}$ in cases where it recurs on itself.

The pointer-refinement of the machine requires parameterizing the $\mathcal{LL}$ function with a store used to resolve pointers to continuations. A nice consequence of this parameterization is that we can re-use $\mathcal{LL}$ for *abstract garbage collection* by supplying it an abstract store for the parameter. Doing so only necessitates extending $\mathcal{LL}$ to the case of sets of storable values:

$$\mathcal{LL}_\sigma(S) = \bigcup_{s \in S} \mathcal{LL}_\sigma(s)$$

The CESK* machine incorporates garbage collection by a transition rule that invokes the GC machine as a subroutine to remove garbage from the store (Figure 16). The garbage collection transition introduces non-determinism to the CESK* machine because it applies to any machine state and thus overlaps with the existing transition rules. The non-determinism is interpreted as leaving the choice of *when* to collect garbage up to the machine.

The abstract CESK* incorporates garbage collection by the *concrete garbage collection transition*, *i.e.*, we re-use the definition in Figure 16 with an abstract store, $\hat{\sigma}$, in place of the concrete one. Consequently, it is easy to verify abstract garbage collection approximates its concrete counterpart.

The CESK* machine may collect garbage at any point in the computation, thus an abstract interpretation must soundly approximate *all possible choices* of when to trigger a collection, which the abstract CESK* machine does correctly. This may be a useful analysis *of* garbage collection, however it fails to be a useful analysis *with* garbage collection: for soundness, the abstracted machine must consider the case in which garbage is never collected, implying no storage is reclaimed to improve precision.

However, we can leverage abstract garbage collection to reduce the state-space explored during analysis and to improve precision and analysis time. This is achieved (again) by considering properties of the *concrete* machine, which abstract directly; in this case, we want the concrete machine to deterministically collect garbage. Determinism of the CESK* machine is restored by defining the transition relation as a non-GC transition (Figure 3) followed by the GC transition (Figure 16). This state-space of this concrete machine is "garbage free" and consequently the state-space of the abstracted machine is "abstract garbage free."

In the concrete semantics, a nice consequence of this property is that although continuations are allocated in the store, they are deallocated as soon as they become unreachable, which corresponds to when they would be popped from the stack in a non-pointer-refined machine. Thus the concrete machine really manages continuations like a stack.

Similarly, in the abstract semantics, continuations are deallocated as soon as they become unreachable, which often corresponds to when they would be popped. We say often, because due to the finiteness of the store, this correspondence cannot always hold. However, this approach gives a good finite approximation to infinitary stack analyses that can always match calls and returns.

6. Abstract stack inspection

In this section, we derive an abstract interpreter for the static analysis of a higher-order language with stack inspection. Following the outline of Section 2 and 3, we start from the tail-recursive CM machine of Clements and Felleisen [3], perform a pointer refinement on continuations, then abstract the semantics by a parameterized bounding of the store.

6.1 The $\lambda_{\mathbf{sec}}$-calculus and stack-inspection

The λ_{sec}-calculus of Pottier, Skalka, and Smith is a call-by-value λ-calculus model of higher-order stack inspection [26]. We present the language as given by Clements and Felleisen [3].

All code is statically annotated with a given set of permissions R, chosen from a fixed set $\mathcal{P}$. A computation whose source code was statically annotated with a permission may *enable* that permission for the dynamic extent of a subcomputation. The subcomputation is privileged so long as it is annotated with the same permission, and every intervening procedure call has likewise been annotated with the privilege.

$$\begin{aligned}
e \in \mathit{Exp} ::= \ldots \mid &\texttt{fail} \mid \texttt{(grant } R\ e\texttt{)} \mid \\
&\texttt{(test } R\ e\ e\texttt{)} \mid \texttt{(frame } R\ e\texttt{)}
\end{aligned}$$

A `fail` expression signals an exception if evaluated; by convention it is used to signal a stack-inspection failure. A `(frame` R e`)` evaluates e as the principal R, representing the permissions conferred on e given its origin. A `(grant` R e`)` expression evaluates as e but with the permissions extended with R enabled. A `(test` R e_0 e_1`)` expression evaluates to e_0 if R is enabled and e_1 otherwise.

A trusted annotator consumes a program and the set of permissions it will operate under and inserts frame expressions around each λ-body and intersects all grant expressions with this set of permissions. We assume all programs have been properly annotated.

Stack inspection can be understood in terms of an $\mathcal{OK}$ predicate on an evaluation contexts and permissions. The predicate determines whether the given permissions are enabled for a subexpression in the hole of the context. The $\mathcal{OK}$ predicate holds whenever the context can be traversed from the hole outwards and, for each permission, find an enabling grant context without first finding a denying frame context.

$$\varsigma \longmapsto_{CM} \varsigma'$$

$$\begin{array}{|l|l|}
\hline
\langle \texttt{fail}, \rho, \sigma, \kappa\rangle & \langle \texttt{fail}, \rho, \sigma, \mathbf{mt}^{\emptyset}\rangle \\
\langle (\texttt{frame}\ R\ e), \rho, \sigma, \kappa\rangle & \langle e, \rho, \sigma, \kappa[R \mapsto \text{deny}]\rangle \\
\langle (\texttt{grant}\ R\ e), \rho, \sigma, \kappa\rangle & \langle e, \rho, \sigma, \kappa[R \mapsto \text{grant}]\rangle \\
\langle (\texttt{test}\ R\ e_0\ e_1), \rho, \sigma, \kappa\rangle & \begin{cases} \langle e_0, \rho, \sigma, \kappa\rangle & \text{if } \mathcal{OK}(R, \kappa), \\ \langle e_1, \rho, \sigma, \kappa\rangle & \text{otherwise} \end{cases} \\
\hline
\end{array}$$

$$\begin{array}{rcl}
\mathcal{OK}(\emptyset, \kappa) & & \\
\mathcal{OK}(R, \mathbf{mt}^m) & \iff & (R \cap m^{-1}(\text{deny}) = \emptyset) \\
\left.\begin{array}{r}\mathcal{OK}(R, \mathbf{fn}^m(v, \rho, \kappa)) \\ \mathcal{OK}(R, \mathbf{ar}^m(e, \rho, \kappa))\end{array}\right\} & \iff & \begin{array}{l}(R \cap m^{-1}(\text{deny}) = \emptyset) \wedge \\ \quad \mathcal{OK}(R \setminus m^{-1}(\text{grant}), \kappa)\end{array}
\end{array}$$

Figure 17. The CM machine and $\mathcal{OK}$ predicate.

6.2 The CM machine

The CM (continuation-marks) machine of Clements and Felleisen is a properly tail-recursive extended CESK machine for interpreting higher-order languages with stack-inspection [3].

In the CM machine, continuations are annotated with *marks* [4], which, for the purposes of stack-inspection, are finite maps from permissions to {deny, grant}:

$$\kappa \quad ::= \quad \mathbf{mt}^m \mid \mathbf{ar}^m(e, \rho, \kappa) \mid \mathbf{fn}^m(v, \rho, \kappa).$$

We write $\kappa[R \mapsto c]$ to mean update the marks on κ to $m[R \mapsto c]$.

The CM machine is defined in Figure 17 (transitions that are straightforward adaptations of the corresponding CESK* transitions to incorporate continuation marks are omitted). It relies on the $\mathcal{OK}$ predicate to determine whether the permissions in R are enabled. The $\mathcal{OK}$ predicate performs the traversal of the context (represented as a continuation) using marks to determine which permissions have been granted or denied.

The semantics of a program is given by the set of reachable states from an initial machine configuration:

$$inj_{CM}(e) = \langle e, \emptyset, [a_0 \mapsto \mathbf{mt}^{\emptyset}], a_0\rangle.$$

6.3 The abstract CM* machine

Store-allocating continuations, time-stamping, and bounding the store yields the transition system given in Figure 18. The notation $\hat{\sigma}(a)[R \mapsto c]$ is used to mean $[R \mapsto c]$ should update *some* continuation in $\hat{\sigma}(a)$, *i.e.*,

$$\hat{\sigma}(a)[R \mapsto c] = \hat{\sigma}[a \mapsto \hat{\sigma}(a) \setminus \{\kappa\} \cup \{\kappa[R \mapsto c]\}],$$

for some $\kappa \in \hat{\sigma}(a)$. It is worth noting that continuation marks are updated, not joined, in the abstract transition system.

The $\widehat{\mathcal{OK}^\star}$ predicate (Figure 18) approximates the pointer refinement of its concrete counterpart $\mathcal{OK}$, which can be understood as tracing a path through the store corresponding to traversing the continuation. The abstract predicate holds whenever there exists such a path in the abstract store that would satisfy the concrete predicate: Consequently, in analyzing $(\texttt{test}\ R\ e_0\ e_1)$, e_0 is reachable only when the analysis can prove the $\mathcal{OK}^\star$ predicate holds on some path through the abstract store.

It is straightforward to define a structural abstraction map and verify the abstract CM* machine is a sound approximation of its concrete counterpart:

Theorem 4 (Soundness of the Abstract CM* Machine)**.** *If $\varsigma \longmapsto_{CM} \varsigma'$ and $\alpha(\varsigma) \sqsubseteq \hat{\varsigma}$, then there exists an abstract state $\hat{\varsigma}'$, such that $\hat{\varsigma} \longmapsto_{\widehat{CM^\star_t}} \hat{\varsigma}'$ and $\alpha(\varsigma') \sqsubseteq \hat{\varsigma}'$.*

$$\hat{\varsigma} \longmapsto_{\widehat{CM^\star}} \hat{\varsigma}'$$

$$\begin{array}{|l|l|}
\hline
\langle \texttt{fail}, \rho, \hat{\sigma}, a\rangle & \langle \texttt{fail}, \rho, \hat{\sigma}, a_{\mathbf{mt}}\rangle \\
\langle (\texttt{frame}\ R\ e), \rho, \hat{\sigma}, a\rangle & \langle e, \rho, \hat{\sigma}(a)[R \mapsto \text{deny}], a\rangle \\
\langle (\texttt{grant}\ R\ e), \rho, \hat{\sigma}, a\rangle & \langle e, \rho, \hat{\sigma}(a)[R \mapsto \text{grant}], a\rangle \\
\langle (\texttt{test}\ R\ e_0\ e_1), \rho, \hat{\sigma}, a\rangle & \begin{cases} \langle e_0, \rho, \hat{\sigma}, a\rangle & \text{if } \widehat{\mathcal{OK}^\star}(R, \hat{\sigma}, a), \\ \langle e_1, \rho, \hat{\sigma}, a\rangle & \text{otherwise.} \end{cases} \\
\hline
\end{array}$$

$$\begin{array}{rcl}
\widehat{\mathcal{OK}^\star}(\emptyset, \hat{\sigma}, a) & & \\
\widehat{\mathcal{OK}^\star}(R, \hat{\sigma}, a) & \iff & (R \cap m^{-1}(\text{deny}) = \emptyset) \\
\text{if } \hat{\sigma}(a) \ni \mathbf{mt}^m & & \\
\widehat{\mathcal{OK}^\star}(R, \hat{\sigma}, a) & \iff & (R \cap m^{-1}(\text{deny}) = \emptyset) \wedge \\
\text{if } \hat{\sigma}(a) \ni \mathbf{fn}^m(v, \rho, b) & & \widehat{\mathcal{OK}^\star}(R \setminus m^{-1}(\text{grant}), \hat{\sigma}, b) \\
\text{or } \hat{\sigma}(a) \ni \mathbf{ar}^m(e, \rho, b) & &
\end{array}$$

Figure 18. The abstract CM* machine.

7. Widening to improve complexity

If implemented naïvely, it takes time exponential in the size of the input program to compute the reachable states of the abstracted machines. Consider the size of the state-space for the abstract time-stamped CESK* machine:

$$\begin{aligned}
&|Exp \times Env \times \widehat{Store} \times Addr \times Time| \\
&= |Exp| \times |Addr|^{|Var|} \times |Storable|^{|Addr|} \times |Addr| \times |Time|.
\end{aligned}$$

Without simplifying any further, we clearly have an exponential number of abstract states.

To reduce complexity, we can employ widening in the form of Shivers's single-threaded store [29]. 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 set of partial states plus a single globally approximating store, *i.e.*, $aval : Exp \to System$, where:

$$System = \mathcal{P}\left(Exp \times Env \times Addr \times Time\right) \times \widehat{Store}.$$

We compute this as a fixed point of a monotonic function, f:

$$\begin{aligned}
&f : System \to System \\
&f(C, \hat{\sigma}) = (C', \hat{\sigma}'') \text{ where} \\
&\quad Q' = \left\{(c', \hat{\sigma}') : c \in C \text{ and } (c, \hat{\sigma}) \longmapsto (c', \hat{\sigma}')\right\} \\
&(c_0, \hat{\sigma}_0) \cong inj(e) \\
&\quad C' = C \cup \left\{c' : (c', _) \in Q'\right\} \cup \{c_0\} \\
&\quad \hat{\sigma}'' = \hat{\sigma} \sqcup \bigsqcup_{(_, \hat{\sigma}') \in Q'} \hat{\sigma}',
\end{aligned}$$

so that $aval(e) = \text{lfp}(f)$. The maximum number of iterations of the function f times the cost of each iteration bounds the complexity of the analysis.

Polynomial complexity for monovariance: It is straightforward to compute the cost of a monovariant (in our framework, a "0CFA-like") analysis with this widening. In a monovariant analysis, environments disappear; a monovariant system-space simplifies to:

$$\begin{aligned}
&System_0 \\
&= \mathcal{P}\left(Exp \times Lab \times Lab_\perp\right) \\
&\quad \times (\overbrace{(Var + Lab)}^{\text{addresses}} \to \overbrace{(Exp \times Lab)}^{\textbf{fn} \text{ conts}} + \overbrace{(Exp \times Lab)}^{\textbf{ar} \text{ conts}} + Lam).
\end{aligned}$$

If ascended monotonically, one could add one new partial state each time or introduce a new entry into the global store. Thus, the maximum number of monovariant iterations is:

$$|Exp| \times |Lab|^2 + 1 \\ + |Var + Lab| \times (2|Exp \times Lab| + |Lam|),$$

which is cubic in the size of the program.

8. Related work

The study of abstract machines for the λ-calculus began with Landin's SECD machine [21], the theory of abstract interpretation with the POPL papers of the Cousots' [6, 7], and static analysis of the λ-calculus with Jones's coupling of abstract machines and abstract interpretation [17]. All three have been active areas of research since their inception, but only recently have well known abstract machines been connected with abstract interpretation by Midtgaard and Jensen [23, 24]. We strengthen the connection by demonstrating a general technique for abstracting abstract machines.

Abstract interpretation of abstract machines: The approximation of abstract machine states for the analysis of higher-order languages goes back to Jones [17], who argued abstractions of regular tree automata could solve the problem of recursive structure in environments. We re-invoked that wisdom to eliminate the recursive structure of continuations by allocating them in the store.

Midtgaard and Jensen present a 0CFA for a CPS λ-calculus language [23]. The approach is based on Cousot-style calculational abstract interpretation [5], applied to a functional language. Like the present work, Midtgaard and Jensen start with an "off-the-shelf" abstract machine for the concrete semantics (in this case, the CE machine of Flanagan, *et al.* [14]) and employ a reachable-states model. They then compose well-known Galois connections to reveal a 0CFA with reachability in the style of Ayers [2].[6] The CE machine is not sufficient to interpret direct-style programs, so the analysis is specialized to programs in continuation-passing style. Later work by Midtgaard and Jensen went on to present a similar calculational abstract interpretation treatment of a monomorphic CFA for an ANF λ-calculus [24]. The concrete semantics are based on reachable states of the $\mathrm{C}_a\mathrm{EK}$ machine [14]. The abstract semantics approximate the control stack component of the machine by its top element, which is similar to the labeled machine abstraction given in Section 2.7 when $k = 0$.

Although our approach is not calculational like Midtgaard and Jensen's, it continues in their tradition by applying abstract interpretation to off-the-shelf tail-recursive machines. We extend the application to direct-style machines for a k-CFA-like abstraction that handles tail calls, laziness, state, exceptions, first-class continuations, and stack inspection. We have extended *return flow analysis* to a completely direct style (no ANF or CPS needed) within a framework that accounts for polyvariance.

Harrison gives an abstract interpretation for a higher-order language with control and state for the purposes of automatic parallelization [15]. Harrison maps Scheme programs into an imperative intermediate language, which is interpreted on a novel abstract machine. The machine uses a procedure string approach similar to that given in Section 2.7 in that the store is addressed by procedure strings. Harrison's first machine employs higher-order values to represent functions and continuations and he notes, "the straightforward abstraction of this semantics leads to abstract domains containing higher-order objects (functions) over reflexive domains, whereas our purpose requires a more concrete compile-time representation of the values assumed by variables. We therefore modify the semantics such that its abstraction results in domains which are both finite and non-reflexive." Because of the reflexivity of denotable values, a direct abstraction is not possible, so he performs closure conversion on the (representation of) the semantic function. Harrison then abstracts the machine by bounding the procedure string space (and hence the store) via an abstraction he calls stack configurations, which is represented by a finite set of members, each of which describes an infinite set of procedure strings.

To prove that Harrison's abstract interpreter is correct he argues that the machine interpreting the translation of a program in the intermediate language corresponds to interpreting the program as written in the standard semantics—in this case, the denotational semantics of $\mathrm{R}^3\mathrm{RS}$. On the other hand, our approach relies on well known machines with well known relations to calculi, reduction semantics, and other machines [10, 8]. These connections, coupled with the strong similarities between our concrete and abstract machines, result in minimal proof obligations in comparison. Moreover, programs are analyzed in direct-style under our approach.

Abstract interpretation of lazy languages: Jones has analyzed non-strict functional languages [17, 16], but that work has only focused on the by-name aspect of laziness and does not address memoization as done here. Sestoft examines flow analysis for lazy languages and uses abstract machines to prove soundness [27]. In particular, Sestoft presents a lazy variant of Krivine's machine similar to that given in Section 3 and proves analysis is sound with respect to the machine. Likewise, Sestoft uses Landin's SECD machine as the operational basis for proving globalization optimizations correct. Sestoft's work differs from ours in that analysis is developed separately from the abstract machines, whereas we derive abstract interpreters directly from machine definitions. Faxén uses a type-based flow analysis approach to analyzing a functional language with explicit thunks and evals, which is intended as the intermediate language for a compiler of a lazy language [9]. In contrast, our approach makes no assumptions about the typing discipline and analyzes source code directly.

Realistic language features and garbage collection: Static analyzers typically hemorrhage precision in the presence of exceptions and first-class continuations: they jump to the top of the lattice of approximation when these features are encountered. Conversion to continuation- and exception-passing style can handle these features without forcing a dramatic ascent of the lattice of approximation [29]. The cost of this conversion, however, is lost knowledge—both approaches obscure static knowledge of stack structure, by desugaring it into syntax.

Might and Shivers introduced the idea of using abstract garbage collection to improve precision and efficiency in flow analysis [25]. They develop a garbage collecting abstract machine for a CPS language and prove it correct. We extend abstract garbage collection to direct-style languages interpreted on the CESK machine.

Static stack inspection: Most work on the static verification of stack inspection has focused on type-based approaches. Skalka and Smith present a type system for static enforcement of stack-inspection [30]. Pottier *et al.* present type systems for enforcing stack-inspection developed via a static correspondence to the dynamic notion of security-passing style [26]. Skalka *et al.* present type and effect systems that use linear temporal logic to express regular properties of program traces and show how to statically enforce both stack- and history-based security mechanisms [31]. Our approach, in contrast, is not typed-based and focuses only on stack-inspection, although it seems plausible the approach of Section 6 extends to the more general history-based mechanisms.

[6] Ayers derived an abstract interpreter by transforming (the representation of) a denotational continuation semantics of Scheme into a state transition system (an abstract machine), which he then approximated using Galois connections [2].

9. Conclusions and perspective

We have demonstrated the utility of store-allocated continuations by deriving novel abstract interpretations of the CEK, a lazy variant of Krivine's, and the stack-inspecting CM machines. These abstract interpreters are obtained by a straightforward pointer refinement and structural abstraction that bounds the address space, making the abstract semantics safe and computable. Our technique allows concrete implementation technology to be mapped straightforwardly into that of static analysis, which we demonstrated by incorporating abstract garbage collection and optimizations to avoid abstract space leaks, both of which are based on existing accounts of concrete GC and space efficiency. Moreover, the abstract interpreters properly model tail-calls by virtue of their concrete counterparts being properly tail-call optimizing. Finally, our technique uniformly scales up to richer language features. We have supported this by extending the abstract CESK machine to analyze conditionals, first-class control, exception handling, and state. We speculate that store-allocating bindings and continuations is sufficient for a straightforward abstraction of most existing machines.

Acknowledgments: We thank Matthias Felleisen, Jan Midtgaard, and Sam Tobin-Hochstadt for discussions and suggestions. We also thank the anonymous reviewers for their close reading and helpful critiques; their comments have improved this paper.

References

[1] Mads S. Ager, Olivier Danvy, and Jan Midtgaard. A functional correspondence between call-by-need evaluators and lazy abstract machines. *Information Processing Letters*, 90(5):223–232, June 2004.

[2] Andrew E. Ayers. *Abstract analysis and optimization of Scheme*. PhD thesis, Massachusetts Institute of Technology, 1993.

[3] John Clements and Matthias Felleisen. A tail-recursive machine with stack inspection. *ACM Trans. Program. Lang. Syst.*, 26(6):1029–1052, November 2004.

[4] John Clements, Matthew Flatt, and Matthias Felleisen. Modeling an algebraic stepper. In *ESOP '01: Proceedings of the 10th European Symposium on Programming Languages and Systems*, pages 320–334, 2001.

[5] Patrick Cousot. The calculational design of a generic abstract interpreter. In M. Broy and R. Steinbrüggen, editors, *Calculational System Design*. 1999.

[6] Patrick Cousot and Radhia Cousot. Abstract interpretation: A unified lattice model for static analysis of programs by construction or approximation of fixpoints. In *Conference Record of the Fourth ACM Symposium on Principles of Programming Languages*, pages 238–252, 1977.

[7] Patrick Cousot and Radhia Cousot. Systematic design of program analysis frameworks. In *POPL '79: Proceedings of the 6th ACM SIGACT-SIGPLAN Symposium on Principles of Programming Languages*, pages 269–282, 1979.

[8] Olivier Danvy. *An Analytical Approach to Program as Data Objects*. DSc thesis, Department of Computer Science, Aarhus University, October 2006.

[9] Karl Faxén. Optimizing lazy functional programs using flow inference. In *Static Analysis*, pages 136–153. 1995.

[10] Matthias Felleisen. *The Calculi of Lambda-v-CS Conversion: A Syntactic Theory of Control and State in Imperative Higher-Order Programming Languages*. PhD thesis, Indiana University, 1987.

[11] Matthias Felleisen, Robert B. Findler, and Matthew Flatt. *Semantics Engineering with PLT Redex*. August 2009.

[12] Matthias Felleisen and Daniel P. Friedman. Control operators, the SECD-machine, and the lambda-calculus. In *3rd Working Conference on the Formal Description of Programming Concepts*, August 1986.

[13] Mattias Felleisen and D. P. Friedman. A calculus for assignments in higher-order languages. In *POPL '87: Proceedings of the 14th ACM SIGACT-SIGPLAN Symposium on Principles of Programming Languages*, pages 314+, 1987.

[14] Cormac Flanagan, Amr Sabry, Bruce F. Duba, and Matthias Felleisen. The essence of compiling with continuations. In *PLDI '93: Proceedings of the ACM SIGPLAN 1993 Conference on Programming Language Design and Implementation*, pages 237–247, June 1993.

[15] Williams L. Harrison. The interprocedural analysis and automatic parallelization of scheme programs. *LISP and Symbolic Computation*, 2(3):179–396, October 1989.

[16] N. Jones and N. Andersen. Flow analysis of lazy higher-order functional programs. *Theoretical Computer Science*, 375(1-3):120–136, May 2007.

[17] Neil D. Jones. Flow analysis of lambda expressions (preliminary version). In *Proceedings of the 8th Colloquium on Automata, Languages and Programming*, pages 114–128, 1981.

[18] Neil D. Jones and Steven S. Muchnick. A flexible approach to interprocedural data flow analysis and programs with recursive data structures. In *POPL '82: Proceedings of the 9th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*, pages 66–74, 1982.

[19] Jean-Louis Krivine. Un interpréteur du lambda-calcul. 1985.

[20] Jean-Louis Krivine. A call-by-name lambda-calculus machine. *Higher-Order and Symbolic Computation*, 20(3):199–207, September 2007.

[21] Peter J. Landin. The mechanical evaluation of expressions. *The Computer Journal*, 6(4):308–320, 1964.

[22] Jan Midtgaard. Control-flow analysis of functional programs. Technical Report BRICS RS-07-18, DAIMI, Department of Computer Science, University of Aarhus, December 2007. To appear in revised form in ACM Computing Surveys.

[23] Jan Midtgaard and Thomas Jensen. A calculational approach to control-flow analysis by abstract interpretation. In María Alpuente and Germán Vidal, editors, *SAS*, volume 5079 of *Lecture Notes in Computer Science*, pages 347–362, 2008.

[24] Jan Midtgaard and Thomas P. Jensen. Control-flow analysis of function calls and returns by abstract interpretation. In *ICFP '09: Proceedings of the 14th ACM SIGPLAN International Conference on Functional Programming*, pages 287–298, 2009.

[25] Matthew Might and Olin Shivers. Improving flow analyses via ΓCFA: Abstract garbage collection and counting. In *ICFP '06: Proceedings of the Eleventh ACM SIGPLAN International Conference on Functional Programming*, pages 13–25, 2006.

[26] François Pottier, Christian Skalka, and Scott Smith. A systematic approach to static access control. *ACM Trans. Program. Lang. Syst.*, 27(2):344–382, March 2005.

[27] Peter Sestoft. *Analysis and efficient implementation of functional programs*. PhD thesis, University of Copenhagen, October 1991.

[28] Zhong Shao and Andrew W. Appel. Space-efficient closure representations. In *LFP '94: Proceedings of the 1994 ACM Conference on LISP and Functional Programming*, pages 150–161, 1994.

[29] Olin G. Shivers. *Control-Flow Analysis of Higher-Order Languages*. PhD thesis, Carnegie Mellon University, 1991.

[30] Christian Skalka and Scott Smith. Static enforcement of security with types. In *ICFP '00: Proceedings of the fifth ACM SIGPLAN International Conference on Functional Programming*, pages 34–45, September 2000.

[31] Christian Skalka, Scott Smith, and David Van Horn. Types and trace effects of higher order programs. *Journal of Functional Programming*, 18(02):179–249, 2008.