Pushdown Flow Analysis of First-Class Control

Dimitrios Vardoulakis Olin Shivers
Northeastern University
dimvar@ccs.neu.edu shivers@ccs.neu.edu

Abstract

Pushdown models are better than control-flow graphs for higher-order flow analysis. They faithfully model the call/return structure of a program, which results in fewer spurious flows and increased precision. However, pushdown models require that calls and returns in the analyzed program nest properly. As a result, they cannot be used to analyze language constructs that break call/return nesting such as generators, coroutines, `call/cc`, *etc*.

In this paper, we extend the CFA2 flow analysis to create the first pushdown flow analysis for languages with first-class control. We modify the abstract semantics of CFA2 to allow continuations to escape to, and be restored from, the heap. We then present a summarization algorithm that handles escaping continuations via a new kind of summary edge. We prove that the algorithm is sound with respect to the abstract semantics.

Categories and Subject Descriptors F.3.2 [*Semantics of Programming Languages*]: Program Analysis

General Terms Languages

Keywords pushdown flow analysis, first-class continuations, restricted continuation-passing style, summarization, higher-order functional language

1. Introduction

Function call and return is the fundamental control-flow mechanism in higher-order languages. Therefore, if a flow analysis is to model program behavior faithfully, it must handle call and return well. Pushdown models of programs [15, 17, 22] enable flow analyses with unbounded call/return matching. These analyses are more precise than analyses based on control-flow graphs.

Pushdown models require that calls and returns in the analyzed program nest properly. However, many control constructs, some of them in mainstream programming languages, break call/return nesting. *Generators* (*e.g.*, in Python [14] or JavaScript [8]) are functions that are usually called inside loops to produce a sequence of values one at a time. A generator executes until it reaches a `yield` statement, at which point it returns the value passed to `yield` to its calling context. When the generator is called again, execution resumes at the first instruction after the `yield`. *Coroutines* (*e.g.*, in Simula67 [3] or Lua [11]) can also suspend and resume their execution, but are more expressive than generators because they can specify where to pass control when they yield. *First-class continuations* reify the rest of the computation as a value. Continuations allow complex control flow, such as jumping back to functions that have already returned. Undelimited continuations (`call/cc` in Scheme [20] and SML/NJ [1]) capture the entire stack. Delimited continuations [4, 7], found in Scala [16] and some Schemes, capture part of the stack. Continuations can express generators and coroutines, as well as multi-threading [18, 25] and Prolog-style backtracking. All these operators provide a rich variety of control behaviors. Unfortunately, we cannot currently use pushdown models to analyze programs that use them.

We rectify this situation by extending the CFA2 flow analysis [22] to languages with first-class control. In this article, we make the following contributions.

- CFA2 is based on abstract interpretation of programs in continuation-passing style (*abbrev.* CPS). We present a CFA2-style abstract semantics for Restricted CPS, a variant of CPS that allows continuations to escape, but also permits effective reasoning about the stack [23]. When we detect a continuation that may escape, we copy the stack into the heap (sec. 4.3). We prove that the abstract semantics is a safe approximation of the actual run-time behavior of the program (sec. 4.4).
- In pushdown flow analysis, each state has a stack of unbounded size. Hence, the state space is infinite. Algorithms that explore the state space use a technique called *summarization*. First-class control causes the stack to be copied into the heap, so our analysis must also deal with infinitely many heaps. We show that it is not necessary to keep continuations in the heap during summarization; we handle escaping continuations using a new kind of summary edge (sec. 5.3).
- When calls and returns nest properly, execution paths satisfy a property called *unique decomposition:* for each state s in a path, we can uniquely identify another state s' as the entry of the procedure that contains s [17]. In the presence of first-class control, a state can belong to more than one procedure. We allow paths that are decomposable in multiple ways and prove that our analysis is sound (sec. 5.4).
- If continuations escape upward, a flow analysis cannot generally avoid including extra, spurious control flows that degrade precision. What about continuations that are only used downward, such as exception handlers or continuations captured by `call/cc` that never escape? We show that CFA2 can avoid spurious control flows for downward continuations (sec. 5.5).

2. Why pushdown models?

Finite-state flow analyses, such as *k*-CFA, approximate programs as graphs of abstract machine states. Each node in such a graph represents a program point plus some amount of abstracted environment and control context. Every path in the graph is considered a possible execution of the program. Thus, executions are strings in a *regular* language.

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'11, September 19–21, 2011, Tokyo, Japan.
Copyright © 2011 ACM 978-1-4503-0865-6/11/09. . . $10.00

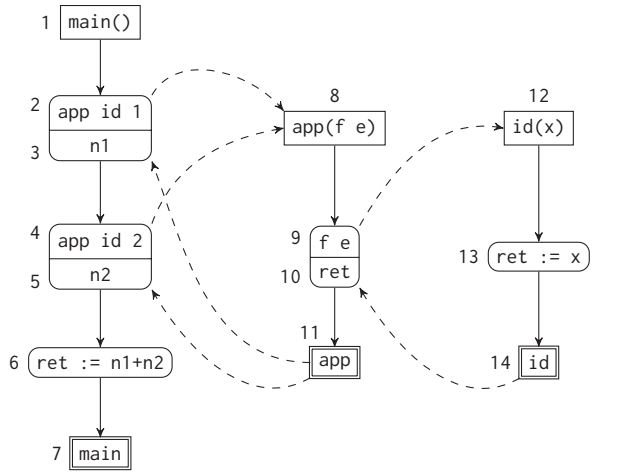

Figure 1. Control-flow graph for a simple program

Finite-state analyses do not handle call and return well. Remembering only a bounded number of pending calls, they must inevitably include spurious paths in which a function is called from one program point and returns to a different one.

Execution traces that match calls with returns are strings in a *context-free* language. Therefore, by abstracting a program to a pushdown automaton (or equivalent), we can use the stack to eliminate call/return mismatch. The following examples illustrate the advantages of pushdown models.

2.1 Data-flow information

The following Scheme program defines the apply and identity functions, then binds `n1` to 1 and `n2` to 2 and adds them. At program point `(+ n1 n2)`, both variables are bound to constants; we would like a static analysis to be able to discover this.

```
(define app (λ (f e) (f e)))
(define id  (λ (x) x))

(let* ((n1 (app id 1))
       (n2 (app id 2)))
  (+ n1 n2))
```

Fig. 1 shows the control-flow graph for this program. In the graph, the top level of the program is presented as a function called `main`. Function entry and exit nodes are rectangles with sharp corners. Inner nodes are rectangles with rounded corners. Each call site is represented by a call node and a corresponding return node, which contains the variable to which the result of the call is assigned. Each function uses a local variable `ret` for its return value. Solid arrows are intraprocedural steps. Dashed arrows go from call sites to function entries and from function exits to return points. There is no edge between call and return nodes; a call reaches its corresponding return only if the callee terminates.

A monovariant analysis, such as 0CFA, considers all paths to be valid executions. Thus, we can bind `n1` to 2 by calling `app` from 4 and returning to 3. Also, we can bind `n2` to 1 by calling `app` from 2 and returning to 5. At point 6, 0CFA thinks that each variable can be bound to either 1 or 2. (For polyvariant analyses, we can create similar examples.) On the other hand, if we only consider paths that respect call/return matching, there is no spurious flow of data. At 6, `n1` and `n2` are bound to constants.

2.2 Stack-change calculation

Besides data-flow information, pushdown models also improve control-flow information. Hence, we can use them to accurately calculate stack changes between program points. With call/return matching, there is only one execution path in our example:

$$1\ 2\overset{8\ 9}{\overset{12\ 13\ 14}{\overset{10\ 11}{3\ 4}}}\overset{8\ 9}{\overset{12\ 13\ 14}{\overset{10\ 11}{5\ 6\ 7}}}$$

In contrast, 0CFA thinks that the program has a loop (there is a path from 4 to 3).

Many optimizations require accurate information about stack change. For instance:

- Most compound data are heap allocated in the general case. Examples include: closure environments, cons pairs, records, objects, *etc*. If we can show statically that such a piece of data is only passed downward, we can allocate it on the stack and reduce garbage-collection overhead.
- Continuations captured by `call/cc` may not escape upward. In this case, we do not need to copy the stack into the heap.
- In object-oriented languages, objects may have methods that are thread-safe by using locks. An escape analysis can eliminate unnecessary locking/unlocking in the methods of thread-private objects.

Such optimizations are better performed with pushdown models.

2.3 Fake rebinding

It is possible that two control-related references to the same variable are always bound in the same run-time environment. If a flow analysis cannot detect this, it will additionally consider execution traces in which the two references are bound to different abstract values. We call this phenomenon fake rebinding [22].

```
(define (compose-same f x) (f (f x)))
```

In `compose-same`, both references to `f` are always bound in the same environment (the top stack frame). However, if multiple closures flow to `f`, a finite-state analysis will consider traces that call one closure at the inner call site and a different closure at the outer call site. CFA2 forbids these paths because it knows that both references are bound in the top frame—it correctly relates *environment* and *control* structure.

2.4 Broadening the applicability of pushdown models

Pushdown-reachability algorithms use a dynamic-programming technique called summarization. Summarization relies on proper nesting of calls and returns. If we call function `app` from location 2 in our example, summarization knows that, if the call returns, it will return to location 3, *not* location 5.

What if the call to `app` is a tail call, in which case the return point is effectively in a different procedure from the call site? We can handle this by creating *cross-procedure* summaries [22].

In languages with exceptions, the return point may be deeper in the stack. We can transform this case into ordinary call/return nesting and handle it precisely with CFA2. Instead of thinking of an exception as a single jump deeper in the stack, we can return to the caller, which checks if it can handle the exception and if not, it passes it to its own caller and so on. Functions return a pair of values, one for normal return and one for exceptional return. The JavaScript implementation of CFA2 [5] uses this technique for exceptions.

But what if the return point has already been popped off the stack, as is the case when using first-class control constructs? Pushdown models cannot currently analyze such programs, so we have to fall back to a finite-state analysis and live with its limitations. In

$$
\begin{array}{rcl}
v \in Var & = & UVar + CVar \\
u \in UVar & = & \text{a set of identifiers} \\
k \in CVar & = & \text{a set of identifiers} \\
\psi \in Lab & = & ULab + CLab \\
l \in ULab & = & \text{a set of labels} \\
\gamma \in CLab & = & \text{a set of labels} \\
lam \in Lam & = & ULam + CLam \\
ulam \in ULam & ::= & (\lambda_l\,(u\,k)\;call) \\
clam \in CLam & ::= & (\lambda_\gamma\,(u)\;call) \\
call \in Call & = & UCall + CCall \\
UCall & ::= & (f\;e\;q)^l \\
CCall & ::= & (q\;e)^\gamma \\
g \in Exp & = & UExp + CExp \\
f, e \in UExp & = & ULam + UVar \\
q \in CExp & = & CLam + CVar \\
pr \in Program & = & ULam
\end{array}
$$

Figure 2. Partitioned CPS

the rest of this paper, we show how to generalize pushdown models to first-class control.

3. Restricted CPS

Preliminary definitions In this section we describe our CPS language. Compilers that use CPS [1, 10, 21] frequently partition the terms in a program into two disjoint sets, the *user* and the *continuation* set, and treat user terms differently from continuation terms. We adopt this partitioning here (Fig. 2). Variables, lambdas and calls get labels from $ULab$ or $CLab$. Labels are pairwise distinct. User lambdas take a user argument and the current continuation; continuation lambdas take only a user argument.

We assume that all variables in a program have distinct names. Then, the defining lambda of a variable v, written $def_\lambda(v)$, is the lambda term that contains v in its list of formals. For any term g, $iu_\lambda(g)$ is the innermost user lambda that contains g. Concrete syntax enclosed in $[\![\cdot]\!]$ denotes an item of abstract syntax. Functions with a '?' subscript are predicates, *e.g.*, $Var_?(e)$ returns true if e is a variable and false otherwise.

We use two notations for tuples, $(e_1, \ldots, e_n)$ and $\langle e_1, \ldots, e_n \rangle$, to avoid confusion when tuples are deeply nested. We use the latter for lists as well; ambiguities will be resolved by the context. Lists are also described by a head-tail notation, *e.g.*, $3 :: \langle 1, 3, -47 \rangle$.

Handling first-class control In CPS, we can naturally express first-class control without using special primitives: when continuations are captured by user closures, they may escape.

Escaping continuations complicate reasoning about the stack. To enable reasoning about the stack in the presence of first-class control, we have previously proposed a syntactically-restricted variant of CPS, called Restricted CPS (*abbrev.* RCPS) [23].

Definition 1 (Restricted CPS)**.** *A program is in Restricted CPS iff a continuation variable can appear free in a user lambda in operator position only.*

In RCPS, continuations escape in a well-behaved way: *after a continuation escapes, it can only be called*; it cannot be passed as an argument again. For example, the CPS-translation of `call/cc`, which is `(λ(f cc) (f (λ(v k) (cc v)) cc))`, is a valid RCPS term. Terms like `(λ(x k) (k (λ(y k2) (y 123 k))))` are not valid.

We can transform this term (and any CPS term) to a valid RCPS term by η-expanding to return the free reference to operator position: `(λ(x k) (k (λ(y k2) (y 123 (λ(u) (k u))))))`. Why do we distinguish these very similar terms? Because, according to the Orbit policy (*cf.* sec. 4.1), their stack behaviors differ. In the case of the first term, when execution reaches `(y 123 k)`, we must restore the environment of the continuation that flows to `k`, which may cause arbitrary change to the stack. In the second case, when execution reaches `(y 123 (λ(u) (k u)))`, a new continuation is born and no stack change is required. Thus, RCPS forces all exotic stack change to happen when calling an escaping continuation, not in other kinds of call sites.

$$
\begin{array}{rcl}
\varsigma \in State & = & Eval + Apply \\
Eval & = & UEval + CEval \\
UEval & = & UCall \times BEnv \times VEnv \times Time \\
CEval & = & CCall \times BEnv \times VEnv \times Time \\
Apply & = & UApply + CApply \\
UApply & = & UClos \times UClos \times CClos \times VEnv \times Time \\
CApply & = & CClos \times UClos \times VEnv \times Time \\
Clos & = & UClos + CClos \\
d \in UClos & = & ULam \times BEnv \\
c \in CClos & = & (CLam \times BEnv) + halt \\
\beta \in BEnv & = & Var \rightharpoonup Time \\
ve \in VEnv & = & Var \times Time \rightharpoonup Clos \\
t \in Time & = & Lab^*
\end{array}
$$

(a) Concrete domains

$$
\mathcal{A}(g, \beta, ve) \triangleq \begin{cases} (g, \beta) & Lam_?(g) \\ ve(g, \beta(g)) & Var_?(g) \end{cases}
$$

$$
\begin{array}{ll}
[\text{UEA}] & ([\![(f\;e\;q)^l]\!], \beta, ve, t) \to (proc, d, c, ve, l :: t) \\
& proc = \mathcal{A}(f, \beta, ve) \\
& d = \mathcal{A}(e, \beta, ve) \\
& c = \mathcal{A}(q, \beta, ve) \\
\\
[\text{UAE}] & (proc, d, c, ve, t) \to (call, \beta', ve', t) \\
& proc \equiv \langle [\![(\lambda_l\,(u\,k)\;call)]\!], \beta \rangle \\
& \beta' = \beta[u \mapsto t][k \mapsto t] \\
& ve' = ve[(u, t) \mapsto d][(k, t) \mapsto c] \\
\\
[\text{CEA}] & ([\![(q\;e)^\gamma]\!], \beta, ve, t) \to (proc, d, ve, \gamma :: t) \\
& proc = \mathcal{A}(q, \beta, ve) \\
& d = \mathcal{A}(e, \beta, ve) \\
\\
[\text{CAE}] & (proc, d, ve, t) \to (call, \beta', ve', t) \\
& proc = \langle [\![(\lambda_\gamma\,(u)\;call)]\!], \beta \rangle \\
& \beta' = \beta[u \mapsto t] \\
& ve' = ve[(u, t) \mapsto d]
\end{array}
$$

(b) Concrete semantics

Figure 3. Concrete semantics and domains

Concrete semantics Execution in RCPS is guided by the semantics of Fig. 3. In the terminology of abstract interpretation [2], this semantics is called the *concrete* semantics. In order to find properties of a program at compile time, one needs to derive a computable approximation of the concrete semantics, called the *abstract* semantics (*cf.* sec. 4).

Execution traces alternate between $Eval$ and $Apply$ states. At an $Eval$ state, we evaluate the subexpressions of a call site before performing a call. At an $Apply$, we perform the call.

The last component of each state is a *time*, which is a sequence of call sites. $Eval$-to-$Apply$ transitions increment the time by recording the label of the corresponding call site. $Apply$-to-$Eval$ transitions leave the time unchanged. Thus, the time t of a state reveals the call sites along the execution path to that state.

Times indicate points in the execution when variables are bound. The binding environment β is a partial function that maps variables to their binding times. The variable environment ve maps variable/time pairs to values. To find the value of a variable v, we look up the time v was put in β, and use that to search for the actual value in ve. By pairing variables with times, we allow a single variable to have multiple bindings at run time.

Let's look at each transition individually. At a $UEval$ state over $[\![(f\ e\ q)^l]\!]$, we use the function $\mathcal{A}$ to evaluate the simple expressions f, e and q: lambdas are paired up with β to become closures, while variables are looked up in ve using β. We add the label l in front of the current time and transition to a $UApply$ state (rule [UEA]).

From $UApply$ to $Eval$, we bind the formals of a procedure $\langle [\![(\lambda_l (u\ k)\ call)]\!], \beta\rangle$ to the arguments and jump to its body. The new binding environment β' extends the procedure's environment, with u and k mapped to the current time. The new variable environment ve' maps (u, t) to the user argument d, and (k, t) to the continuation c (rule [UAE]).

The remaining two transitions are similar. We use $halt$ to denote the top-level continuation of a program pr. The initial state $\mathcal{I}(pr)$ is $((pr, \emptyset), \mathit{input}, halt, \emptyset, \langle\rangle)$, where *input* is a closure of the form $\langle [\![(\lambda_l (u\ k)\ call)]\!], \emptyset\rangle$. The initial time is the empty sequence of call sites.

4. The CFA2 abstraction

In this section, we'll extend the abstract semantics of CFA2 to handle first-class control. The semantics uses two binding environments, a stack and a heap; we also use the stack for return-point information.

We show the actual transition rules in section 4.3; the main difference from the previous semantics is that continuations can now be copied to, and restored from, the heap. Before that, we discuss how to manage the stack in RCPS (sec. 4.1) and how to decide whether a variable reference will be looked up in the stack or the heap (sec. 4.2). In section 4.4, we prove that the abstract semantics is an approximation of the concrete semantics.

The CFA2 abstraction only takes us halfway to a computable analysis: the abstract state space is infinite, so we cannot explore it by enumerating all states. We tackle this problem in section 5.

4.1 Stack-management policy

The Orbit compiler [9, 10] compiles a CPS intermediate representation to final code that uses a stack. Orbit views continuations as closures whose environment record is a stack frame. To decide when to push and pop the stack, we follow Orbit's policy. The main idea behind Orbit's policy is that *we can manage the stack for a CPS program in the same way that we would manage it for the original direct-style program:*

- For every call to a user function, we push a frame for the arguments.
- We pop a frame at function returns. In CPS, user functions "return" by calling the current continuation with a return value.
- We also pop a frame at tail calls. A $UCall$ call site is a tail call in CPS iff it was a tail call in the original direct-style program. In tail calls, the continuation argument is a variable.
- When a continuation is captured by a user closure, we copy the stack into the heap.
- When we call a continuation that has escaped, we restore its stack from the heap.

4.2 Stack/heap split

The stack in CFA2 is more than a control structure for return-point information; it is also an *environment* structure—it contains bindings. CFA2 has a novel approach to variable binding: *two references to the same variable need not be looked up in the same binding environment.* We split references into two categories: stack and heap references. In direct-style, if a reference appears at the same nesting level as its binder, then it is a stack reference, otherwise it is a heap reference. For example, $(\lambda_1\texttt{(x)}(\lambda_2\texttt{(y)(x (x y))))}$ has a stack reference to y and two heap references to x.

Intuitively, only heap references may escape. When we call a user function, we push a frame for its arguments, so we know that stack references are always bound in the top frame. When control reaches a heap reference, its frame is either deeper in the stack, or it has been popped. We look up stack references in the top frame, and heap references in the heap. Stack lookups below the top frame never happen (Fig. 4b).

When a program p is CPS-converted to a program p', stack (*resp.* heap) references in p remain stack (*resp.* heap) references in p'. All references added by the transform are stack references.

We can give an equivalent definition of stack and heap references directly in CPS, without referring to the original direct-style program. Labels can be split into disjoint sets according to the innermost user lambda that contains them. For the CPS translation of the previous program,

$$(\lambda_1\texttt{(x k1)}\ \texttt{(k1}\ (\lambda_2\texttt{(y k2)}\ \texttt{(x y}\ (\lambda_3\texttt{(u)}\ \texttt{(x u k2)}^4))^5))^6)$$

these sets are $\{1, 6\}$ and $\{2, 3, 4, 5\}$. The "label to variable" map $LV(\psi)$ returns all the variables bound by any lambdas that belong in the same set as ψ, *e.g.*, $LV(4) = \{\texttt{y}, \texttt{k2}, \texttt{u}\}$ and $LV(6) = \{\texttt{x}, \texttt{k1}\}$. We use this map to model stack behavior, because all continuation lambdas that "belong" to a given user lambda λ_l get closed by extending λ_l's stack frame (*cf.* section 4.3). Notice that, for any ψ, $LV(\psi)$ contains exactly one continuation variable. Using LV, we give the following definition.

Definition 2 (Stack and heap references)**.**

- *Let ψ be a call site that refers to a variable v. The predicate $S_?(\psi, v)$ holds iff $v \in LV(\psi)$. We call v a* ***stack reference****.*
- *Let ψ be a call site that refers to a variable v. The predicate $H_?(\psi, v)$ holds iff $v \notin LV(\psi)$. We call v a* ***heap reference****.*
- *v is a* ***stack variable****, written $S_?(v)$, iff all its references satisfy $S_?$.*
- *v is a* ***heap variable****, written $H_?(v)$, iff some of its references satisfy $H_?$.*

For instance, $S_?(5, \texttt{y})$ holds because $\texttt{y} \in \{\texttt{y}, \texttt{k2}, \texttt{u}\}$ and $H_?(5, \texttt{x})$ holds because $\texttt{x} \notin \{\texttt{y}, \texttt{k2}, \texttt{u}\}$.

4.3 Abstract semantics

The CFA2 semantics is an abstract machine that executes a program in RCPS (Fig. 4). The abstract domains appear in Fig. 4a. An abstract user closure (member of the set $\widehat{UClos}$) is a set of user lambdas. An abstract continuation closure (member of $\widehat{CClos}$) is either a continuation lambda or $halt$. A frame is a map from variables to abstract values, and a stack is a sequence of frames. All stack operations except $push$ are defined for non-empty stacks only. A heap is a map from variables to abstract values. In contrast to the previous semantics of CFA2, the heap can contain continuation bindings.

Fig. 4c shows the transition rules. First-class control shows up in two of the rules, $[\widehat{\text{UAE}}]$ and $[\widehat{\text{CEA}}]$.

On transition from a $\widehat{UEval}$ state to a $\widehat{UApply}$ state (rule $[\widehat{\text{UEA}}]$), we first evaluate f, e and q. We evaluate atomic user

$$\hat{\varsigma} \in \widehat{UEval} = UCall \times Stack \times Heap$$
$$\hat{\varsigma} \in \widehat{UApply} = ULam \times \widehat{UClos} \times \widehat{CClos} \times Stack \times Heap$$
$$\hat{\varsigma} \in \widehat{CEval} = CCall \times Stack \times Heap$$
$$\hat{\varsigma} \in \widehat{CApply} = \widehat{CClos} \times \widehat{UClos} \times Stack \times Heap$$
$$\hat{d} \in \widehat{UClos} = Pow(ULam)$$
$$\hat{c} \in \widehat{CClos} = CLam + halt$$
$$fr, tf \in Frame = (UVar \rightharpoonup \widehat{UClos}) + (CVar \rightharpoonup \widehat{CClos})$$
$$st \in Stack = Frame^*$$
$$h \in Heap = (UVar \rightharpoonup \widehat{UClos}) + (CVar \rightharpoonup Pow(\widehat{CClos} \times Stack))$$

(a) Abstract domains

$$pop(tf :: st) \triangleq st$$
$$push(fr, st) \triangleq fr :: st$$
$$(tf :: st)(v) \triangleq tf(v)$$
$$(tf :: st)[u \mapsto \hat{d}] \triangleq tf[u \mapsto \hat{d}] :: st$$

(b) Stack operations

$$\hat{\mathcal{A}}_u(e, \psi, st, h) \triangleq \begin{cases} \{e\} & Lam_?(e) \\ st(e) & S_?(\psi, e) \\ h(e) & H_?(\psi, e) \end{cases}$$

$$[\widehat{\text{UEA}}]\ (\llbracket (f\ e\ q)^l \rrbracket, st, h) \rightsquigarrow (ulam, \hat{d}, \hat{c}, st', h)$$
$$ulam \in \hat{\mathcal{A}}_u(f, l, st, h)$$
$$\hat{d} = \hat{\mathcal{A}}_u(e, l, st, h)$$
$$\hat{c} = \begin{cases} st(q) & Var_?(q) \\ q & Lam_?(q) \end{cases}$$
$$st' = \begin{cases} pop(st) & Var_?(q) \\ st & Lam_?(q) \wedge (H_?(l, f) \vee Lam_?(f)) \\ st[f \mapsto \{ulam\}] & Lam_?(q) \wedge S_?(l, f) \end{cases}$$

$$[\widehat{\text{UAE}}]\ (\llbracket (\lambda_l (u\, k)\ call) \rrbracket, \hat{d}, \hat{c}, st, h) \rightsquigarrow (call, st', h')$$
$$st' = push([u \mapsto \hat{d}][k \mapsto \hat{c}], st)$$
$$h'(v) = \begin{cases} h(u) \cup \hat{d} & (v = u) \wedge H_?(u) \\ h(k) \cup \{(\hat{c}, st)\} & (v = k) \wedge H_?(k) \\ h(v) & o/w \end{cases}$$

$$[\widehat{\text{CEA}}]\ (\llbracket (q\ e)^\gamma \rrbracket, st, h) \rightsquigarrow (\hat{c}, \hat{d}, st', h)$$
$$\hat{d} = \hat{\mathcal{A}}_u(e, \gamma, st, h)$$
$$(\hat{c}, st') \in \begin{cases} \{(q, st)\} & Lam_?(q) \\ \{(st(q), pop(st))\} & S_?(\gamma, q) \\ h(q) & H_?(\gamma, q) \end{cases}$$

$$[\widehat{\text{CAE}}]\ (\llbracket (\lambda_\gamma (u)\ call) \rrbracket, \hat{d}, st, h) \rightsquigarrow (call, st', h')$$
$$st' = st[u \mapsto \hat{d}]$$
$$h'(v) = \begin{cases} h(u) \cup \hat{d} & (v = u) \wedge H_?(u) \\ h(v) & o/w \end{cases}$$

(c) Abstract semantics

Figure 4. Abstract semantics and relevant definitions

terms using $\hat{\mathcal{A}}_u$. We non-deterministically choose one of the lambdas that flow to f as the operator in the $\widehat{UApply}$ state.[1] The change to the stack depends on q and f. If q is a variable, the call is a tail call so we pop the stack (case 1). If q is a lambda, it evaluates to a new closure whose environment is the top frame, hence we do not pop the stack (cases 2, 3). Moreover, if f is a lambda or a heap reference then we leave the stack unchanged. However, if f is a stack reference, we set f's value in the top frame to $\{ulam\}$, possibly forgetting other lambdas that flow to f. The strong update to the stack prevents fake rebinding for stack references (*cf.* sec. 2.3): when we return to $\hat{c}$, we may reach more stack references of f. These references and the current one are bound at the same time. Therefore, they must also be bound to $ulam$.

In the $\widehat{UApply}$-to-$\widehat{Eval}$ transition (rule $[\widehat{\text{UAE}}]$), we push a frame for the procedure's arguments. If u is a heap variable, we update its binding in the heap with all lambdas in $\hat{d}$. If k is a heap variable, we have a possibly escaping continuation. We save $\hat{c}$ in the heap and also copy the stack, so that we can restore it if $\hat{c}$ gets called later.

In a $\widehat{CEval}$-to-$\widehat{CApply}$ transition (rule $[\widehat{\text{CEA}}]$), we are preparing for a call to a continuation so we must reset the stack to the stack of its birth. When q is a lambda, it is a newly created closure so the stack does not change. When q is a stack reference, the $\widehat{CEval}$ state is a function return and the continuation's environment is the second stack frame. Therefore, we pop a frame before calling $\hat{c}$. When q is a heap reference, we are calling a continuation that may have escaped. The stack change since the continuation capture can be arbitrary. We non-deterministically pick a pair $(\hat{c}, st')$ from $h(q)$, jump to $\hat{c}$ and restore st', which contains bindings for the stack references in $\hat{c}$.

In the $\widehat{CApply}$-to-$\widehat{Eval}$ transition (rule $[\widehat{\text{CAE}}]$), the top frame is the environment of $\llbracket (\lambda_\gamma (u)\ call) \rrbracket$; stack references in $call$ need this frame on the top of the stack. Hence, we do not push; we extend the top frame with the binding for the continuation's parameter u. If u is a heap variable, we also update the heap.

Example Let's see how the abstract semantics works on a program with `call/cc`. Consider the program

```
(call/cc (λ(c) (somefun (c 42))))
```

where *somefun* is an arbitrary function. We use `call/cc` to capture the top-level continuation and bind it to `c`. Then, *somefun* will never be called, because `(c 42)` will return to the top level with 42 as the result.

The CPS translation of `call/cc` is

```
(λ1(f cc) (f (λ2(x k2) (cc x)) cc))
```

The CPS translation of its argument is

```
(λ3(c k) (c 42 (λ4(u) (somefun_CPS u k))))
```

The initial state $\hat{\mathcal{I}}(pr) = (\lambda_1,\ \{\lambda_3\},\ halt,\ \langle\rangle,\ \emptyset)$ is a $\widehat{UApply}$. (We abbreviate lambdas by their labels.) We push a frame and jump to the body of λ_1. Since `cc` is a heap variable, we save the continuation and the stack in the heap, producing a heap h with a single element $[\texttt{cc} \mapsto \{(halt, \langle\rangle)\}]$, and $\widehat{UEval}$ state

$$(\llbracket (\texttt{f}\ \lambda_2\ \texttt{cc}) \rrbracket,\ \langle [\texttt{f} \mapsto \{\lambda_3\}][\texttt{cc} \mapsto halt] \rangle,\ h).$$

λ_2 is essentially a continuation reified as a user value. We tail call to λ_3, so we pop the stack, producing $\widehat{UApply}$ state

$$(\lambda_3,\ \{\lambda_2\},\ halt,\ \langle\rangle,\ h).$$

[1] An abstract execution explores one path, but the algorithm that searches the state space considers all possible executions.

$$|([\![(g_1 \ldots g_n)^{\psi}]\!], \beta, ve, t)|_{ca} = ([\![(g_1 \ldots g_n)^{\psi}]\!],\ toStack(LV(\psi), \beta, ve),\ |ve|_{ca})$$

$$|(\langle [\![(\lambda_l\,(u\,k)\ call)]\!], \beta\rangle, d, c, ve, t)|_{ca} = ([\![(\lambda_l\,(u\,k)\ call)]\!], |d|_{ca}, |c|_{ca}, st, |ve|_{ca})$$
$$\text{where}\ \ st = \begin{cases} \langle\rangle & c = halt \\ toStack(LV(\gamma), \beta', ve) & c = ([\![(\lambda_\gamma\,(u')call')]\!], \beta') \end{cases}$$

$$|(\langle [\![(\lambda_\gamma\,(u)\ call)]\!], \beta\rangle, d, ve, t)|_{ca} = ([\![(\lambda_\gamma\,(u)\ call)]\!],\ |d|_{ca}, toStack(LV(\gamma), \beta, ve),\ |ve|_{ca})$$

$$|(halt, d, ve, t)|_{ca} = (halt, |d|_{ca},\ \langle\rangle,\ |ve|_{ca})$$

$$|([\![(\lambda_l\,(u\,k)\ call)]\!], \beta)|_{ca} = \{[\![(\lambda_l\,(u\,k)\ call)]\!]\}$$

$$|([\![(\lambda_\gamma\,(u)\ call)]\!], \beta)|_{ca} = [\![(\lambda_\gamma\,(u)\ call)]\!]$$

$$|halt|_{ca} = halt$$

$$|ve|_{ca} = \{\,(u, \textstyle\bigcup_t |ve(u,t)|_{ca})\ :\ (u \in UVar) \wedge H_?(u)\} \cup \{\,(k, \bigcup_t makecs(ve(k,t), ve))\ :\ (k \in CVar) \wedge H_?(k)\}$$
$$\text{where}\ \ makecs(c, ve) \triangleq \begin{cases} (halt, \langle\rangle) & c = halt \\ ([\![(\lambda_\gamma\,(u')call)]\!], toStack(LV(\gamma), \beta, ve)) & c = \langle [\![(\lambda_\gamma\,(u')call)]\!], \beta\rangle \end{cases}$$

$$toStack(\{u_1, \ldots, u_n, k\}, \beta, ve) \triangleq \begin{cases} \langle [\,\overline{u_i \mapsto \hat{d}_i}\,][k \mapsto halt]\rangle & ve(k, \beta(k)) = halt \\ [\,\overline{u_i \mapsto \hat{d}_i}\,][k \mapsto [\![(\lambda_\gamma\,(u)\ call)]\!]] :: st & ve(k, \beta(k)) = ([\![(\lambda_\gamma\,(u)\ call)]\!], \beta') \end{cases}$$
$$\text{where}\ \ \hat{d}_i = |ve(u_i, \beta(u_i))|_{ca}\ \ \text{and}\ \ st = toStack(LV(\gamma), \beta', ve)$$

Figure 5. From concrete states to abstract states

$$g \sqsubseteq g \quad \text{where} \quad g \in (halt + Lam + Call)$$

$$\langle a_1, \ldots, a_n\rangle \sqsubseteq \langle b_1, \ldots, b_n\rangle \quad \text{iff} \quad \text{for } 1 \leqslant i \leqslant n,\ a_i \sqsubseteq b_i$$

$$\hat{d}_1 \sqsubseteq \hat{d}_2 \quad \text{iff} \quad \hat{d}_1 \subseteq \hat{d}_2$$

$$h_1 \sqsubseteq h_2 \quad \text{iff} \quad h_1(v) \sqsubseteq h_2(v) \quad \text{for each } v \in \mathrm{dom}(h_1)$$

$$tf_1 :: st_1 \sqsubseteq tf_2 :: st_2 \quad \text{iff} \quad tf_1 \sqsubseteq tf_2\ \wedge\ st_1 \sqsubseteq st_2$$

$$\langle\rangle \sqsubseteq \langle\rangle$$

$$tf_1 \sqsubseteq tf_2 \quad \text{iff} \quad tf_1(v) \sqsubseteq tf_2(v) \quad \text{for each } v \in \mathrm{dom}(tf_1)$$

Figure 6. The $\sqsubseteq$ relation on abstract states

We next push a frame and jump to the body of λ_3:

$$([\![(\mathtt{c}\ \ 42\ \ \lambda_4)]\!],\ \langle[\mathtt{c} \mapsto \{\lambda_2\}][k \mapsto halt]\rangle,\ h).$$

This is a non-tail call, so we do not pop:

$$(\lambda_2,\ \{42\},\ \lambda_4,\ \langle[\mathtt{c} \mapsto \{\lambda_2\}][\mathtt{k} \mapsto halt]\rangle,\ h).$$

We push a frame and jump to the body of λ_2:

$$([\![(\mathtt{cc}\ \ \mathtt{x})]\!],\ \langle[\mathtt{x} \mapsto \{42\}][\mathtt{k2} \mapsto \lambda_4],\ [\mathtt{c} \mapsto \{\lambda_2\}][\mathtt{k} \mapsto halt]\rangle,\ h).$$

As cc is a heap reference, we ignore the current continuation and stack and restore $(halt, \langle\rangle)$ from the heap:

$$(halt,\ \{42\},\ \langle\rangle,\ h).$$

The program terminates with value $\{42\}$.

4.4 Correctness of the abstract semantics

In this section, we show that the abstract semantics *simulates* the concrete semantics, which means that the execution of a program under the abstract semantics is a safe approximation of its actual run-time behavior. First, we define a map $|\cdot|_{ca}$ from concrete to abstract states. Next, we show that if ς transitions to ς' in the concrete semantics, the abstract counterpart $|\varsigma|_{ca}$ of ς transitions to a state $\hat{\varsigma}'$ which approximates $|\varsigma'|_{ca}$. Therefore, each concrete execution, *i.e.*, sequence of states related by $\rightarrow$, has a corresponding abstract execution that computes an approximate answer.

The map $|\cdot|_{ca}$ appears in Fig. 5. The abstraction of an $Eval$ state ς of the form $([\![(g_1 \ldots g_n)^{\psi}]\!], \beta, ve, t)$ is an $\widehat{Eval}$ state $\hat{\varsigma}$ with the same call site. Since ς does not have a stack, we must expose stack-related information hidden in β and ve. Assume that λ_l is the innermost user lambda that contains ψ. To reach ψ, control passed from a $\widehat{UApply}$ state $\hat{\varsigma}'$ over λ_l. According to our stack policy, the top frame contains bindings for the formals of λ_l and any temporaries added along the path from $\hat{\varsigma}'$ to $\hat{\varsigma}$. Therefore, the domain of the top frame is a subset of $LV(l)$, *i.e.*, a subset of $LV(\psi)$. For each user variable $u_i \in (LV(\psi) \cap \mathrm{dom}(\beta))$, the top frame contains $[u_i \mapsto |ve(u_i, \beta(u_i))|_{ca}]$. Let k be the sole continuation variable in $LV(\psi)$. If $ve(k, \beta(k))$ is $halt$ (the return continuation is the top-level continuation), the rest of the stack is empty. If $ve(k, \beta(k))$ is $([\![(\lambda_\gamma\,(u)\ call)]\!], \beta')$, the second frame is for the user lambda in which λ_γ was born, and so forth: proceeding through the stack, we add a frame for each live activation of a user lambda until we reach $halt$.

The abstraction of a $UApply$ state over $\langle [\![(\lambda_l\,(u\,k)\ call)]\!], \beta\rangle$ is a $\widehat{UApply}$ state $\hat{\varsigma}$ whose operator is $[\![(\lambda_l\,(u\,k)\ call)]\!]$. The stack of $\hat{\varsigma}$ represents the environment in which the continuation argument was created, and we compute it using $toStack$ as above.

Abstracting a $CApply$ is similar to the $UApply$ case, only now the top frame is the environment of the continuation operator. Note that the abstraction maps drop the time of the concrete states, since the abstract states do not use times.

The abstraction of a user closure is the singleton set with the corresponding lambda. The abstraction of a continuation closure is the corresponding lambda.

The abstraction $|ve|_{ca}$ of a variable environment is a heap, which contains bindings for the user and the continuation heap variables. Each heap user variable is bound to the set of lambdas of the closures that can flow to it. Each heap continuation variable k is bound to a set of continuation-stack pairs. For each closure that can flow to k, we create a pair with the lambda of that closure and the corresponding stack.

The relation $\hat{\varsigma}_1 \sqsubseteq \hat{\varsigma}_2$ is a partial order on abstract states and can be read as "$\hat{\varsigma}_1$ is more precise than $\hat{\varsigma}_2$" (Fig. 6). Tuples are ordered pointwise. Abstract user closures are ordered by inclusion. Two stacks are in $\sqsubseteq$ iff they have the same length and the corresponding frames are in $\sqsubseteq$.

We can now state the simulation theorem. The proof proceeds by case analysis on the concrete transition relation; full details can be found in the first author's forthcoming dissertation.

Theorem 3 (Simulation)**.** *If $\varsigma \to \varsigma'$ and $|\varsigma|_{ca} \sqsubseteq \hat{\varsigma}$, then there exists $\hat{\varsigma}'$ such that $\hat{\varsigma} \rightsquigarrow \hat{\varsigma}'$ and $|\varsigma'|_{ca} \sqsubseteq \hat{\varsigma}'$.*

5. Exploring the infinite state space

Pushdown-reachability algorithms, including CFA2, deal with the unbounded stack size by using a dynamic-programming technique called *summarization*. These algorithms work on transition systems whose stack is unbounded, but the rest of the components are bounded. Due to escaping continuations, we also have to deal with infinitely many heaps.

5.1 Overview of summarization

We start with an informal overview of summarization. Assume that a program is executing and control reaches the entry of a procedure. We start computing inside the procedure. While doing so, we are visiting several program points inside the procedure and possibly calling (and returning from) other procedures. Sometime later, we reach the exit and are about to return to the caller with a result. The intuition behind summarization is that, during this computation, the return point was irrelevant; it influences reachability only after we return to the caller. Consequently, if from a program point n with an empty stack we can reach a point n' with stack s', then from n with stack s we can reach n' with stack $\mathit{append}(s', s)$.

Let's use summarization to find which nodes of the graph of Fig. 1 are reachable from node 1. We find reachable nodes by recording *path edges*, *i.e.*, edges whose source is the entry of a procedure and target is some program point in the same procedure. Path edges should not be confused with the edges already present in the graph. They are artificial edges used by the analysis to represent intraprocedural paths, hence the name.

Node 1 goes to 2, so we record the edges $\langle 1, 1\rangle$ and $\langle 1, 2\rangle$. From 2 we call `app`, so we record the call $\langle 2, 8\rangle$ and jump to 8. In `app`, we find path edges $\langle 8, 8\rangle$ and $\langle 8, 9\rangle$. We find a new call $\langle 9, 12\rangle$ and jump to 12. Inside `id`, we discover the edges $\langle 12, 12\rangle$, $\langle 12, 13\rangle$ and $\langle 12, 14\rangle$. Edges that go from an entry to an exit, such as $\langle 12, 14\rangle$, are called *summary* edges. We have not been keeping track of the stack, so we use the recorded calls to find the return point. The only call to `id` is $\langle 9, 12\rangle$, so 14 returns to 10 and we find a new edge $\langle 8, 10\rangle$, which leads to $\langle 8, 11\rangle$. We record $\langle 8, 11\rangle$ as a summary also. From the call $\langle 2, 8\rangle$, we see that 11 returns to 3, so we record edges $\langle 1, 3\rangle$ and $\langle 1, 4\rangle$. Now, we have a new call to `app`. Reachability inside `app` does not depend on its calling context. From the summary $\langle 8, 11\rangle$, we know that 4 can reach 5, so we find $\langle 1, 5\rangle$. Subsequently, we find the last two path edges, which are $\langle 1, 6\rangle$ and $\langle 1, 7\rangle$.

During the search, we did two kinds of transitions. The first kind includes intraprocedural steps and calls; these transitions do not shrink the stack. The second is function returns, which shrink the stack. Since we are not keeping track of the stack, we find the target nodes of the second kind of transitions in an indirect way, by recording calls and summaries. We show a summarization-based algorithm for CFA2 in section 5.3. The next section describes the local semantics, which we use in the algorithm for transitions that do not shrink the stack.

$$
\begin{aligned}
\widehat{\mathit{Eval}} &= \mathit{Call} \times \widehat{\mathit{Stack}} \times \widehat{\mathit{Heap}} \\
\widehat{\mathit{UApply}} &= \mathit{ULam} \times \widehat{\mathit{UClos}} \times \widehat{\mathit{Heap}} \\
\widehat{\mathit{CApply}} &= \widehat{\mathit{CClos}} \times \widehat{\mathit{UClos}} \times \widehat{\mathit{Stack}} \times \widehat{\mathit{Heap}} \\
\widehat{\mathit{Frame}} &= \mathit{UVar} \rightharpoonup \widehat{\mathit{UClos}} \\
\widehat{\mathit{Stack}} &= \widehat{\mathit{Frame}} \\
\widehat{\mathit{Heap}} &= \mathit{UVar} \rightharpoonup \widehat{\mathit{UClos}}
\end{aligned}
$$

(a) Local domains

$$
\begin{aligned}
&|(\mathit{call}, \mathit{st}, h)|_{al} = (\mathit{call}, |\mathit{st}|_{al}, |h|_{al}) \\
&|(\mathit{ulam}, \hat{d}, \hat{c}, \mathit{st}, h)|_{al} = (\mathit{ulam}, \hat{d}, |h|_{al}) \\
&|(\hat{c}, \hat{d}, \mathit{st}, h)|_{al} = (\hat{c}, \hat{d}, |\mathit{st}|_{al}, |h|_{al}) \\
&|\mathit{st}|_{al} = \begin{cases} \emptyset & \mathit{st} = \langle\rangle \\ \mathit{tf} \upharpoonright \mathit{UVar} & \mathit{st} = \mathit{tf} :: \mathit{st}' \end{cases} \\
&|h|_{al} = h \upharpoonright \mathit{UVar}
\end{aligned}
$$

(b) Abstract to local maps

$$
\tilde{\mathcal{A}}_u(e, \psi, \mathit{tf}, h) \triangleq \begin{cases} \{e\} & \mathit{Lam}_?(e) \\ \mathit{tf}(e) & \mathit{S}_?(\psi, e) \\ h(e) & \mathit{H}_?(\psi, e) \end{cases}
$$

$$
\begin{aligned}
&[\widehat{\mathrm{UEA}}] \quad (\llbracket (f\ e\ q)^l \rrbracket, \mathit{tf}, h) \twoheadrightarrow (\mathit{ulam}, \hat{d}, h) \\
&\qquad \mathit{ulam} \in \tilde{\mathcal{A}}_u(f, l, \mathit{tf}, h) \\
&\qquad \hat{d} = \tilde{\mathcal{A}}_u(e, l, \mathit{tf}, h)
\end{aligned}
$$

$$
\begin{aligned}
&[\widehat{\mathrm{UAE}}] \quad (\llbracket (\lambda_l\, (u\, k)\ \mathit{call}) \rrbracket, \hat{d}, h) \twoheadrightarrow (\mathit{call}, [u \mapsto \hat{d}], h') \\
&\qquad h'(v) = \begin{cases} h(u) \cup \hat{d} & (v = u) \wedge H_?(u) \\ h(v) & o/w \end{cases}
\end{aligned}
$$

$$
\begin{aligned}
&[\widehat{\mathrm{CEA}}] \quad (\llbracket (\mathit{clam}\ e)^\gamma \rrbracket, \mathit{tf}, h) \twoheadrightarrow (\mathit{clam}, \hat{d}, \mathit{tf}, h) \\
&\qquad \hat{d} = \tilde{\mathcal{A}}_u(e, \gamma, \mathit{tf}, h)
\end{aligned}
$$

$$
\begin{aligned}
&[\widehat{\mathrm{CAE}}] \quad (\llbracket (\lambda_\gamma\, (u)\ \mathit{call}) \rrbracket, \hat{d}, \mathit{tf}, h) \twoheadrightarrow (\mathit{call}, \mathit{tf}', h') \\
&\qquad \mathit{tf}' = \mathit{tf}[u \mapsto \hat{d}] \\
&\qquad h'(v) = \begin{cases} h(u) \cup \hat{d} & (v = u) \wedge H_?(u) \\ h(v) & o/w \end{cases}
\end{aligned}
$$

(c) Local semantics

Figure 7. Local semantics and relevant definitions

5.2 Local semantics

Summarization operates on a finite set of program points. Since the abstract state space is infinite, we cannot use abstract states as program points. For this reason, we introduce *local states* (Fig. 7a) and define a map $|\cdot|_{al}$ from abstract to local states (Fig. 7b).

The local semantics (Fig. 7) describes executions that do not touch the rest of the stack (*i.e.*, executions where functions do not return). A $\widetilde{CEval}$ state with call site $[\![(k\ e)^\gamma]\!]$ has no successor in this semantics. Since functions do not call their continuations, the local frames and heaps contain only user bindings. Local steps are otherwise similar to abstract steps. Note that there is no provision for first-class control in the local transitions; they are identical to the previous ones [22]. The metavariable $\tilde{\varsigma}$ ranges over local states. We define the map $|\cdot|_{cl}$ from concrete to local states to be $|\cdot|_{al} \circ |\cdot|_{ca}$.

Summarization distinguishes between different kinds of states: entries, exits, calls, returns and inner states. CPS lends itself naturally to such a categorization. The following definition works for all three state spaces: concrete, abstract and local.

Definition 4 (Classification of states)**.**

- *A UApply state is an **Entry**—control is about to enter the body of a function.*
- *A CEval state is an **Exit-Ret** when the operator is a stack reference—a function is about to return a result to its context.*
- *A CEval state is an **Exit-Esc** when the operator is a heap reference—we are calling a continuation that may have escaped.*
- *A CEval state where the operator is a lambda is an **Inner** state.*
- *A UEval state is an **Exit-TC** when the continuation argument is a variable—at tail calls control does not return to the caller.*
- *A UEval state is a **Call** when the continuation argument is a lambda.*
- *A CApply state is a **Return** if its predecessor is an exit, or an **Inner** state if its predecessor is also an inner state. Our algorithm does not distinguish between the two kinds of CApplys; the difference is just conceptual.*

5.3 The CFA2 algorithm

The algorithm for CFA2 appears in Fig. 8. Its goal is to compute which local states are reachable from the initial state of a program.

For readers familiar with the previous algorithm: the main addition is the handling of Exit-Esc states (lines 25-33). Escaping continuations also require changes to the handling of entries and tail calls. Entries are now a separate case, instead of together with $\widetilde{CApply}$s and inner $\widetilde{CEval}$s. Last, the `Propagate` function takes an extra argument.

Structure of the algorithm The algorithm uses a workset W, which contains path edges and summaries to be examined. An edge $(\tilde{\varsigma}_1, \tilde{\varsigma}_2)$ is an ordered pair of local states. We call $\tilde{\varsigma}_1$ the *source* and $\tilde{\varsigma}_2$ the *target* of the edge. At every iteration, we remove an edge from W and process it, potentially adding new edges in W. We stop when W is empty.

An edge $(\tilde{\varsigma}_1, \tilde{\varsigma}_2)$ is a summary when $\tilde{\varsigma}_1$ is an entry and $\tilde{\varsigma}_2$ is either an Exit-Ret or an Exit-Esc, not necessarily in the same procedure. Summaries carry an important message: *each continuation that can be passed to $\tilde{\varsigma}_1$ can flow to the operator position of $\tilde{\varsigma}_2$.*

The algorithm maintains several sets. The results of the analysis are stored in the set $Seen$. It contains path edges (from a procedure entry to a state in the same procedure) and summary edges. The target of an edge in $Seen$ is reachable from the source and from the initial state. Summaries are also stored in $Summary$. $Escapes$ contains Exit-Esc states. If the continuation parameter of a user lambda is a heap variable, entries over that lambda are stored in $EntriesEsc$. $Final$ records final states, *i.e.*, $\widetilde{CApply}$s that call $halt$ with a return value for the whole program. $Callers$ contains triples $\langle \tilde{\varsigma}_1, \tilde{\varsigma}_2, \tilde{\varsigma}_3 \rangle$, where $\tilde{\varsigma}_1$ is an entry, $\tilde{\varsigma}_2$ is a call in the same procedure and $\tilde{\varsigma}_3$ is the entry of the callee. $TCallers$ contains triples $\langle \tilde{\varsigma}_1, \tilde{\varsigma}_2, \tilde{\varsigma}_3 \rangle$, where $\tilde{\varsigma}_1$ is an entry, $\tilde{\varsigma}_2$ is a tail call in the same procedure and $\tilde{\varsigma}_3$ is the entry of the callee. The initial state $\tilde{\mathcal{I}}(pr)$ is defined as $|\mathcal{I}(pr)|_{cl}$. The helper function $succ(\tilde{\varsigma})$ returns the successor(s) of $\tilde{\varsigma}$ according to the local semantics.

Summaries for first-class continuations Perhaps surprisingly, even though continuations can escape to the heap in the abstract semantics, we do not need continuations in the local heap. We can handle escaping continuations with summaries. Consider the example from section 4.3. When control reaches $[\![(\texttt{cc x})]\!]$, we want to find which continuation flows to `cc`. We know that $def_\lambda(\texttt{cc})$ is λ_1. By looking at the single $\widehat{UApply}$ over λ_1, we find that $halt$ flows to `cc`. This suggests that, for escaping continuations, we need summaries of the form $(\tilde{\varsigma}_1, \tilde{\varsigma}_2)$ where $\tilde{\varsigma}_2$ is an Exit-Esc over a call site $[\![(k\ e)^\gamma]\!]$ and $\tilde{\varsigma}_1$ is an entry over $def_\lambda(k)$.

Edge processing Each edge $(\tilde{\varsigma}_1, \tilde{\varsigma}_2)$ is processed in one of six ways, depending on $\tilde{\varsigma}_2$. If $\tilde{\varsigma}_2$ is a return or an inner state (line 12), then its successor $\tilde{\varsigma}_3$ is a state in the same procedure. Since $\tilde{\varsigma}_2$ is reachable from $\tilde{\varsigma}_1$, $\tilde{\varsigma}_3$ is also reachable from $\tilde{\varsigma}_1$. If we have not already recorded the edge $(\tilde{\varsigma}_1, \tilde{\varsigma}_3)$, we do it now (line 44).

If $\tilde{\varsigma}_2$ is a call (line 14) then $\tilde{\varsigma}_3$ is the entry of the callee, so we propagate $(\tilde{\varsigma}_3, \tilde{\varsigma}_3)$ instead of $(\tilde{\varsigma}_1, \tilde{\varsigma}_3)$ (line 16). Also, we record the call in $Callers$. If an exit $\tilde{\varsigma}_4$ is reachable from $\tilde{\varsigma}_3$, it should return to the continuation born at $\tilde{\varsigma}_2$ (line 18). The function `Update` is responsible for computing the return state. We find the return value $\hat{d}$ by evaluating the expression e_4 passed to the continuation (lines 48-49). Since we are returning to λ_{γ_2}, we must restore the environment of its creation, which is tf_2 (possibly with stack filtering, line 50). The new state $\tilde{\varsigma}$ is the corresponding return of $\tilde{\varsigma}_2$, so we propagate $(\tilde{\varsigma}_1, \tilde{\varsigma})$ (lines 51-52).

If $\tilde{\varsigma}_2$ is an Exit-Ret and $\tilde{\varsigma}_1$ is the initial state (lines 19-20), then $\tilde{\varsigma}_2$'s successor is a final state (lines 53-54). If $\tilde{\varsigma}_1$ is some other entry, we record the edge in $Summary$ and pass the result of $\tilde{\varsigma}_2$ to the callers of $\tilde{\varsigma}_1$ (lines 22-23). Last, consider the case of a tail call $\tilde{\varsigma}_4$ to $\tilde{\varsigma}_1$ (line 24). No continuation is born at $\tilde{\varsigma}_4$. Thus, we must find where $\tilde{\varsigma}_3$ (the entry that led to the tail call) was called from. Then again, all calls to $\tilde{\varsigma}_3$ may be tail calls, in which case we keep searching further back in the call chain to find a return point. We do the backward search by transitively adding a cross-procedure summary from $\tilde{\varsigma}_3$ to $\tilde{\varsigma}_2$.

Let $\tilde{\varsigma}_2$ be an Exit-Esc over a call site $[\![(k\ e)^\gamma]\!]$ (line 25). Its predecessor $\tilde{\varsigma}'$ is an entry or a $\widetilde{CApply}$. To reach $\tilde{\varsigma}_2$, the algorithm must go through $\tilde{\varsigma}'$. Hence, the first time the algorithm sees $\tilde{\varsigma}_2$ is at line 7 or 13, which means that $\tilde{\varsigma}_1$ is an entry over $iu_\lambda([\![(k\ e)^\gamma]\!])$ and $(\tilde{\varsigma}_1, \tilde{\varsigma}_2)$ is not in $Summary$. Thus, the test at line 26 is true. We record $\tilde{\varsigma}_2$ in $Escapes$. We also create summaries from entries over $def_\lambda(k)$ to $\tilde{\varsigma}_2$, in order to find which continuations can flow to k. We make sure to put these summaries in $Summary$ (line 29), so that when they are examined, the test at line 26 is false.

When $\tilde{\varsigma}_2$ is examined again, this time $(\tilde{\varsigma}_1, \tilde{\varsigma}_2)$ is in $Summary$. If $\tilde{\varsigma}_1$ is the initial state, $\tilde{\varsigma}_2$ can call $halt$ and transition to a final state (line 30). Otherwise, we look for calls to $\tilde{\varsigma}_1$ to find continuations that can be called at $\tilde{\varsigma}_2$ (line 32). If there are tail calls to $\tilde{\varsigma}_1$, we propagate summaries transitively (line 33).

If $\tilde{\varsigma}_2$ is an entry over $[\![(\lambda_l(u\,k)\ call)]\!]$, its successor $\tilde{\varsigma}_3$ is a state in the same procedure, so we propagate $(\tilde{\varsigma}_1, \tilde{\varsigma}_3)$ (lines 6-7). If k is a heap variable (lines 8-9), we put $\tilde{\varsigma}_2$ in $EntriesEsc$ (so that it can be found from line 29). Also, if we have seen Exit-Esc states that call k, we create summaries from $\tilde{\varsigma}_2$ to those states (line 11).

If $\tilde{\varsigma}_2$ is a tail call (line 34), we find its successors and record the call in $TCallers$ (lines 35-37). If a successor of $\tilde{\varsigma}_2$ goes to an exit, we propagate a cross-procedure summary transitively (line 41). Moreover, if $\tilde{\varsigma}_4$ is an Exit-Esc, we want to make sure that $(\tilde{\varsigma}_1, \tilde{\varsigma}_4)$ is in $Summary$ when it is examined. We cannot call `Propagate` with `true` at line 41 because we would be mutating $Summary$ while iterating over it. Instead, we use a temporary set which we unite with $Summary$ after the loop (line 42).

```
01  Summary, Callers, TCallers, EntriesEsc, Escapes, Final ⟵ ∅
02  Seen, W ⟵ {(Ĩ(pr), Ĩ(pr))}
03  while W ≠ ∅
04    remove (ς̃1, ς̃2) from W
05    switch ς̃2
06      case ς̃2 of Entry
07        for the ς̃3 in succ(ς̃2), Propagate(ς̃1, ς̃3, false)
08        ς̃2 of the form ([[(λl(u k) call)]], d̂, h)
09        if H?(k) then
10          insert ς̃2 in EntriesEsc
11          for each ς̃3 in Escapes that calls k, Propagate(ς̃2, ς̃3, true)
12      case ς̃2 of CApply, Inner-CEval
13        for the ς̃3 in succ(ς̃2), Propagate(ς̃1, ς̃3, false)
14      case ς̃2 of Call
15        for each ς̃3 in succ(ς̃2)
16          Propagate(ς̃3, ς̃3, false)
17          insert (ς̃1, ς̃2, ς̃3) in Callers
18          for each (ς̃3, ς̃4) in Summary, Update(ς̃1, ς̃2, ς̃3, ς̃4)
19      case ς̃2 of Exit-Ret
20        if ς̃1 = Ĩ(pr) then Final(ς̃2)
21        else
22          insert (ς̃1, ς̃2) in Summary
23          for each (ς̃3, ς̃4, ς̃1) in Callers, Update(ς̃3, ς̃4, ς̃1, ς̃2)
24          for each (ς̃3, ς̃4, ς̃1) in TCallers, Propagate(ς̃3, ς̃2, false)
25      case ς̃2 of Exit-Esc
26        if (ς̃1, ς̃2) not in Summary then
27          insert ς̃2 in Escapes
28          ς̃2 of the form ([[(k e)^γ]], tf, h)
29          for each ς̃3 in EntriesEsc over def_λ(k), Propagate(ς̃3, ς̃2, true)
30        else if ς̃1 = Ĩ(pr) then Final(ς̃2)
31        else
32          for each (ς̃3, ς̃4, ς̃1) in Callers, Update(ς̃3, ς̃4, ς̃1, ς̃2)
33          for each (ς̃3, ς̃4, ς̃1) in TCallers, Propagate(ς̃3, ς̃2, true)
34      case ς̃2 of Exit-TC
35        for each ς̃3 in succ(ς̃2)
36          Propagate(ς̃3, ς̃3, false)
37          insert (ς̃1, ς̃2, ς̃3) in TCallers
38          S ⟵ ∅
39          for each (ς̃3, ς̃4) in Summary
40            insert (ς̃1, ς̃4) in S
41            Propagate(ς̃1, ς̃4, false)
42          Summary ⟵ Summary ∪ S

    Propagate(ς̃1, ς̃2, esc) ≜
43    if esc then insert (ς̃1, ς̃2) in Summary
44    if (ς̃1, ς̃2) not in Seen then insert (ς̃1, ς̃2) in Seen and W

    Update(ς̃1, ς̃2, ς̃3, ς̃4) ≜
45    ς̃1 of the form ([[(λ_{l1}(u1 k1) call1)]], d̂1, h1)
46    ς̃2 of the form ([[(f e2 (λ_{γ2} (u2) call2))^{l2}]], tf2, h2)
47    ς̃3 of the form ([[(λ_{l3}(u3 k3) call3)]], d̂3, h2)
48    ς̃4 of the form ([[(k4 e4)^{γ4}]], tf4, h4)
49    d̂ ⟵ Ã_u(e4, γ4, tf4, h4)
50    tf ⟵ { tf2[f ↦ {[[(λ_{l3}(u3 k3) call3)]]}]   S?(l2, f)
           { tf2                                    H?(l2, f) ∨ Lam?(f)
51    ς̃ ⟵ ([[(λ_{γ2}(u2) call2)]], d̂, tf, h4)
52    Propagate(ς̃1, ς̃, false)

    Final(ς̃) ≜
53    ς̃ of the form ([[(k e)^γ]], tf, h)
54    insert (halt, Ã_u(e, γ, tf, h), ∅, h) in Final
```

Figure 8. CFA2 workset algorithm

5.4 Soundness

The local state space is finite, so there are finitely many path and summary edges. We record edges as seen when we insert them in W, which ensures that no edge is inserted in W twice. Therefore, the algorithm always terminates.

We obviously cannot visit an infinite number of abstract states. To establish soundness, we relate the results of the algorithm to the abstract semantics: we show that if a state $\hat{\varsigma}$ is reachable from $\hat{\mathcal{I}}(pr)$, then the algorithm visits $|\hat{\varsigma}|_{al}$ (*cf.* theorem 8).

First-class continuations create an intricate call/return structure, which complicates reasoning about soundness. When calls and returns nest properly, an execution path can be decomposed so that for each state $\hat{\varsigma}$, we can uniquely identify another state $\hat{\varsigma}'$ as the entry of the procedure that contains $\hat{\varsigma}$ [17]. When we add tail calls into the mix, unique decomposition is still possible [24].

However, in the presence of first-class control, a state can belong to *more than one* procedure. For instance, suppose we want to find the entry of the procedure containing $\hat{\varsigma}$ in the following path

$$\hat{\mathcal{I}}(pr) \leadsto^* \hat{\varsigma}_c \leadsto \hat{\varsigma}_e \leadsto^* \hat{\varsigma}'_c \leadsto \hat{\varsigma}'_e \leadsto^* \hat{\varsigma}' \leadsto \hat{\varsigma}$$

where $\hat{\varsigma}'$ is an Exit-Esc over $[\![(k\ e)^\gamma]\!]$, $\hat{\varsigma}_e$ and $\hat{\varsigma}'_e$ are entries over $def_\lambda(k)$, $\hat{\varsigma}_c$ and $\hat{\varsigma}'_c$ are calls. The two entries have the form

$$\hat{\varsigma}_e = (def_\lambda(k), \hat{d}, \hat{c}, st, h)$$
$$\hat{\varsigma}'_e = (def_\lambda(k), \hat{d}', \hat{c}', st', h')$$

Both $\hat{c}$ and $\hat{c}'$ can flow to k and we can call either at $\hat{\varsigma}'$. If we choose to restore $\hat{c}$ and st, then $\hat{\varsigma}$ is in the same procedure as $\hat{\varsigma}_c$. If we restore $\hat{c}'$ and st', $\hat{\varsigma}$ is in the same procedure as $\hat{\varsigma}'_c$. However, it is possible that $\hat{c} = \hat{c}'$ and $st = st'$, in which case $\hat{\varsigma}$ belongs to two procedures. Unique decomposition no longer holds.

For this reason, we now define a set of corresponding entries for each state, instead of a single entry [22].

Definition 5 (Corresponding Entries)**.**
Let $p \equiv \hat{\varsigma}_e \leadsto^* \hat{\varsigma}$ where $\hat{\varsigma}_e$ is an entry. We define $CE_p(\hat{\varsigma})$ to be the smallest set such that:

- *if $\hat{\varsigma}$ is an entry, $CE_p(\hat{\varsigma}) = \{\hat{\varsigma}\}$*
- *if $p \equiv \hat{\varsigma}_e \leadsto^* \hat{\varsigma}_1 \leadsto^+ \hat{\varsigma}$, $\hat{\varsigma}$ is an Exit-Esc over $[\![(k\ e)^\gamma]\!]$, $\hat{\varsigma}_1$ is an entry over $def_\lambda(k)$, then $\hat{\varsigma}_1 \in CE_p(\hat{\varsigma})$.*
- *if $p \equiv \hat{\varsigma}_e \leadsto^* \hat{\varsigma}_1 \leadsto \hat{\varsigma}$, $\hat{\varsigma}$ is neither an entry nor an Exit-Esc, $\hat{\varsigma}_1$ is neither an Exit-Ret nor an Exit-Esc, then $CE_p(\hat{\varsigma}) = CE_p(\hat{\varsigma}_1)$.*
- *if $p \equiv \hat{\varsigma}_e \leadsto^* \hat{\varsigma}_1 \leadsto \hat{\varsigma}_2 \leadsto^* \hat{\varsigma}_3 \leadsto^+ \hat{\varsigma}_4 \leadsto \hat{\varsigma}$, $\hat{\varsigma}$ is a $\widehat{CApply}$ of the form $(\hat{c}, \hat{d}, st, h)$, $\hat{\varsigma}_4$ is an Exit-Esc, $\hat{\varsigma}_3 \in CE_p(\hat{\varsigma}_4)$ and has the form $(ulam, \hat{d}', \hat{c}, st, h')$, $\hat{\varsigma}_2 \in CE^*_p(\hat{\varsigma}_3)$, $\hat{\varsigma}_1$ is a Call, then $CE_p(\hat{\varsigma}_1) \subseteq CE_p(\hat{\varsigma})$.*
- *if $p \equiv \hat{\varsigma}_e \leadsto^* \hat{\varsigma}_1 \leadsto \hat{\varsigma}_2 \leadsto^+ \hat{\varsigma}_3 \leadsto \hat{\varsigma}$, $\hat{\varsigma}$ is a $\widehat{CApply}$, $\hat{\varsigma}_3$ is an Exit-Ret, $\hat{\varsigma}_2 \in CE^*_p(\hat{\varsigma}_3)$, $\hat{\varsigma}_1$ is a Call, then $CE_p(\hat{\varsigma}_1) \subseteq CE_p(\hat{\varsigma})$.*

For each state $\hat{\varsigma}$, we also define $CE^*_p(\hat{\varsigma})$ to be the set of entries that can reach an entry in $CE_p(\hat{\varsigma})$ through tail calls.

Definition 6. *Let $p \equiv \hat{\varsigma}_e \leadsto^* \hat{\varsigma}$ where $\hat{\varsigma}_e$ is an entry. We define $CE^*_p(\hat{\varsigma})$ to be the smallest set such that:*

- *$CE_p(\hat{\varsigma}) \subseteq CE^*_p(\hat{\varsigma})$*
- *if $p \equiv \hat{\varsigma}_e \leadsto^* \hat{\varsigma}_1 \leadsto \hat{\varsigma}_2 \leadsto^* \hat{\varsigma}$, $\hat{\varsigma}_2 \in CE_p(\hat{\varsigma})$, $\hat{\varsigma}_1$ is a Tail Call, then $CE^*_p(\hat{\varsigma}_1) \subseteq CE^*_p(\hat{\varsigma})$.*

Note that if $\hat{\varsigma}$ is an Exit-Esc over $[\![(k\ e)^\gamma]\!]$, a procedure that contains $\hat{\varsigma}$ has an entry $\hat{\varsigma}'$ over $iu_\lambda([\![(k\ e)^\gamma]\!])$. Thus, $\hat{\varsigma}'$ is not in $CE_p(\hat{\varsigma})$ because $iu_\lambda([\![(k\ e)^\gamma]\!]) \neq def_\lambda(k)$. For all other states, $CE_p(\hat{\varsigma})$ is the set of entries of procedures that contain $\hat{\varsigma}$. The following lemma relates the stack of a state with the stacks of its corresponding entries.

Lemma 7. *Let $p \equiv \hat{\mathcal{I}}(pr) \leadsto^* \hat{\varsigma}$ where $\hat{\varsigma} \equiv (\ldots, st, h)$.*

1. *If $\hat{\varsigma}$ is a final state then $CE_p(\hat{\varsigma}) = \emptyset$.*
2. *If $\hat{\varsigma}$ is an entry then $CE_p(\hat{\varsigma}) \neq \emptyset$. (Thus, $CE^*_p(\hat{\varsigma}) \neq \emptyset$.)*
   Let $\hat{\varsigma}_e \in CE^*_p(\hat{\varsigma})$, of the form $(ulam, \hat{d}, \hat{c}, st_e, h_e)$. Then, $st = st_e$ and the continuation argument of $\hat{\varsigma}$ is $\hat{c}$.
3. *If $\hat{\varsigma}$ is an Exit-Esc then its stack is not empty and $CE_p(\hat{\varsigma}) \neq \emptyset$. (We do not assert anything about the stack change between a state in $CE^*_p(\hat{\varsigma})$ and $\hat{\varsigma}$; it can be arbitrary.)*
4. *If $\hat{\varsigma}$ is none of the above, then $CE_p(\hat{\varsigma}) \neq \emptyset$.*
   Let $\hat{\varsigma}_e = ([\![(\lambda_l\,(u\ k)\ call)]\!], \hat{d}, \hat{c}, st_e, h_e)$.
   If $\hat{\varsigma}_e \in (CE^*_p(\hat{\varsigma}) \setminus CE_p(\hat{\varsigma}))$ then
   - *there is a frame tf such that $st \equiv tf :: st_e$, and*
   - *there is a variable k' such that $tf(k') = \hat{c}$.*

   If $\hat{\varsigma}_e \in CE_p(\hat{\varsigma})$ then there is a frame tf such that $st \equiv tf :: st_e$, $\mathrm{dom}(tf) \subseteq LV(l)$, $tf(u) \sqsubseteq \hat{d}$, and $tf(k) = \hat{c}$.

The proof of lemma 7 proceeds by induction on the length of the path p. We now state the soundness theorem. Its proof and the proof of lemma 7 can be found in the first author's forthcoming dissertation.

Theorem 8 (Soundness)**.**
If $p \equiv \hat{\mathcal{I}}(pr) \leadsto^* \hat{\varsigma}$ then, after summarization:

- *If $\hat{\varsigma}$ is a final state, then $|\hat{\varsigma}|_{al} \in Final$.*
- *If $\hat{\varsigma}$ is not final and $\hat{\varsigma}' \in CE_p(\hat{\varsigma})$, then $(|\hat{\varsigma}'|_{al}, |\hat{\varsigma}|_{al}) \in Seen$.*
- *If $\hat{\varsigma}$ is an Exit-Ret or Exit-Esc and $\hat{\varsigma}' \in CE^*_p(\hat{\varsigma})$, then $(|\hat{\varsigma}'|_{al}, |\hat{\varsigma}|_{al}) \in Seen$.*

CFA2 without first-class control is complete, which means that there is no loss in precision when going from abstract to local states [24]. The algorithm of Fig. 8 is not complete; it may compute flows that never happen in the abstract semantics. Consider the code:

```
(define esc (λ(f cc) (f (λ(x k) (cc x)) cc)))

(esc (λ1(v1 k1) (v1 "foo" k1))
     (λ(a) (halt a)))

(esc (λ2(v2 k2) (k2 "bar"))
     (λ(b) (halt b)))
```

In this program, `esc` is the CPS translation of `call/cc`. The two user functions λ_1 and λ_2 expect a reified continuation as their first argument; λ_1 uses that continuation and λ_2 does not. The abstract semantics finds that {`"foo"`} flows to `a` and {`"bar"`} flows to `b`.

However, the workset algorithm thinks that {`"foo"`,`"bar"`} flows to `b`. At the second call to `esc`, it connects the entry to the Exit-Esc state over $[\![$`(cc x)`$]\!]$ at line 11, which is a spurious flow.

5.5 Various approaches to downward continuations

In RCPS, the general form of a user lambda that binds a heap continuation variable is

$$(\lambda_1(\mathtt{u\ k})\ (\ldots(\lambda_2(\mathtt{u2\ k2})\ (\ldots(\mathtt{k\ e})^\gamma\ \ldots))\ \ldots))$$

where λ_1 contains a user lambda λ_2, which in turn contains a heap reference to `k` in operator position.

During execution, if a closure over λ_2 escapes upward, merging of continuations at $[\![$`(k e)`$^\gamma]\!]$ is unavoidable. However, when λ_2 is not passed upward, the abstract semantics still merges at $[\![$`(k e)`$^\gamma]\!]$. A natural question to ask is how precise can CFA2 be for downward continuations, such as exception handlers or continuations captured by `call/cc` that never escape. In both cases, we can avoid merging.

In section 2.4, we described how the JavaScript implementation of CFA2 handles exception throws precisely. Another way to achieve this is by uniformly passing two continuations to each user function, the current continuation and an exception handler [1]. Consider a user lambda $[\![(\lambda(\texttt{u k1 k2})\ (\ldots(\texttt{k2 e})^{\gamma}\ \ldots))]\!]$ where $S_?(\gamma, \texttt{k2})$ holds. Every Exit-Ret over $[\![(\texttt{k2 e})^{\gamma}]\!]$ is an exception throw. The handler continuation lives somewhere on the stack. To find it, we propagate transitive summaries for calls, as we do for tail calls. When the algorithm finds an edge $(\tilde{\varsigma}_1, \tilde{\varsigma}_2)$ where $\tilde{\varsigma}_2$ is an Exit-Ret over $[\![(\texttt{k2 e})^{\gamma}]\!]$, it searches in *Callers* for a triple $(\tilde{\varsigma}_3, \tilde{\varsigma}_4, \tilde{\varsigma}_1)$. If the second continuation argument of $\tilde{\varsigma}_4$ is a lambda, we have found a handler. If not, we propagate a summary $(\tilde{\varsigma}_3, \tilde{\varsigma}_2)$, which has the effect of looking for a handler deeper in the stack. Note that the algorithm must keep these new summaries separate from the other summaries, so as not to confuse exceptional with ordinary control flow.

For continuations captured by `call/cc` that are only used downward, we can avoid merging by combining flow analysis and escape analysis. Consider the lambda at the beginning of this subsection. During flow analysis, we track if any closure over λ_2 escapes upward. We do this by checking for summaries $(\tilde{\varsigma}_1, \tilde{\varsigma}_2)$, where $\tilde{\varsigma}_1$ is an entry over λ_1. If λ_2 is contained in a binding reachable from $\tilde{\varsigma}_2$ [12, sec. 4.4.2], then λ_2 is passed upward and we use the heap to look up k at $[\![(\texttt{k e})^{\gamma}]\!]$. Otherwise, we can assume that λ_2 does not escape. Hence, when we see an edge $(\tilde{\varsigma}_1, \tilde{\varsigma}_2)$ where $\tilde{\varsigma}_1$ is an entry over λ_2 and $\tilde{\varsigma}_2$ is an Exit-Esc over $[\![(\texttt{k e})^{\gamma}]\!]$, we treat it as an exception throw. We use the new transitive summaries to search deeper in the stack for a live activation of λ_1, which tells us what flows to k.

6. Related work

The CFA2 workset algorithm is influenced by the functional approach of Sharir and Pnueli [17] and the tabulation algorithm of Reps *et al.* [15]. CFA2 extends these algorithms to first-class functions, introduces the stack/heap split and applies to control constructs that break call/return nesting. Traditional summary edges describe intraprocedural entry-to-exit flows. We have created several kinds of cross-procedure summaries for the various control patterns. Summaries for tail calls describe flows that do not grow the stack. Summaries for exceptions describe flows that grow the stack; the source of the summary may be deeper in the stack than the target. Finally, summaries for first-class control describe flows with arbitrary stack change. The four different kinds of summaries can be conceptually unified because they serve a common purpose: *they connect a continuation passed to a user function with the state that calls it.*

Earl *et al.* have proposed a pushdown higher-order flow analysis that does not use frames [6]. Instead, their analysis allocates all bindings in the heap with context, in the style of *k*-CFA. For $k = 0$, their analysis runs in time $O(n^6)$, where n is the size of the program. Like all pushdown-reachability algorithms, Earl *et al.*'s analysis records pairs of states $(\varsigma_1, \varsigma_2)$ where ς_2 is same-context reachable from ς_1. However, their algorithm does not classify states as entries, exits, calls, *etc*. This has two drawbacks compared to the tabulation algorithm. First, they do not distinguish between path and summary edges. Thus, they have to search the whole set of edges when they look for return points, even though only summaries can contribute to the search. More importantly, path edges are only a small subset of the set S of all edges between same-context reachable states. By not classifying states, their algorithm maintains the whole set S, not just the path edges. In other words, it records edges whose source is not an entry. In the graph of Fig. 1, some of these edges are $\langle 2, 3\rangle$, $\langle 2, 6\rangle$, $\langle 5, 7\rangle$. Such edges slow down the analysis and do not contribute to call/return matching, because they cannot evolve into summary edges. In CFA2, it is possible to disable the use of frames by classifying each reference as a heap reference. The resulting analysis has similar precision to Earl *et al.*'s analysis for $k = 0$. We conjecture that this variant is not a viable alternative in practice, because of the significant loss in precision.

While there is extensive literature on finite-state higher-order flow analysis, little progress has been made in taming the power of `call/cc` and general continuations. Might and Shivers's ΔCFA [12, 13] introduced a notion of "frame strings" to track stack motion; these strings provide a notational vocabulary for describing and distinguishing various sorts of control transfer: recursive call, tail call, return, primitive application, as well as the more exotic control acts that are constructed with first-class control operators. However, the expressiveness of this device is brought low by its eventual regular-expression-based abstraction. Once abstracted, it loses much of its ability to reason about unusual patterns of control flow. We suspect that the infinite-state analytic framework provided by CFA2 could be the missing piece that would enable a ΔCFA-based analysis to be computed without requiring precision-destroying abstractions.

Shivers and Might have also shown how functional coroutines can be constructed with continuations, and then exploited to fuse pipelines of online transducers together into efficient, optimized code [19]. Being able to apply the power of pushdown models such as CFA2 to the transducer-fusion task raises interesting new possibilities. For example, suppose we had a coroutine generator with a recursive control structure—one that walks a binary tree producing the elements at the leaves. We wish to connect this tree-walking generator to a simple iterative coroutine that adds up all the items it receives. Is a pushdown flow analysis powerful enough to fuse the composition of these two coroutines into a single, recursive tree traversal, instead of an awkward, heavyweight implementation that ping-pongs back and forth between two independent stacks?

7. Conclusions

In this paper, we generalize the CFA2 flow analysis to first-class control. We propose an abstract semantics that allows stacks to be copied to the heap, and a summarization algorithm that handles the infinitely many heaps with a new kind of summary edges. With these additions, CFA2 becomes the first pushdown model that analyzes first-class control constructs. Moreover, CFA2 can now analyze the same language features as *k*-CFA, and do it more accurately. Thus, implementors of higher-order languages can use CFA2 as a drop-in replacement of *k*-CFA.

We also revisit the idea of path decomposition to accommodate states that belong to multiple procedures and prove our analysis sound. We show a program for which the abstract semantics gives a different result from the local semantics and conclude that our new summarization algorithm is not complete. We are not certain that first-class control unavoidably leads to incompleteness; we plan to investigate if changes to the algorithm can make it complete. However, it is possible that the abstract semantics describes a machine strictly more expressive than pushdown systems, and that reachability for this machine is not decidable.

Acknowledgements

This article reports on work supported by the Mozilla Foundation, and by the Defense Advanced Research Projects Agency under Air Force Research Laboratory (AFRL/Rome) Contract No. FA8650-10-C-7090 and Cooperative Agreement No. FA8750-10-2-0233. The views expressed are those of the authors and do not reflect the official policy or position of the Department of Defense or the U.S. Government.

References

[1] A. W. Appel. *Compiling with Continuations*. Cambridge University Press, 1992.

[2] P. Cousot and R. Cousot. Abstract interpretation: a unified lattice model for static analysis of programs by construction or approximation of fixpoints. In *Proceedings of the Fourth Annual ACM Symposium on Principles of Programming Languages*, pages 238–252, Jan. 1977.

[3] O.-J. Dahl and K. Nygaard. SIMULA—an ALGOL-based simulation language. *Communications of the ACM*, 9(9):671–678, 1966.

[4] O. Danvy and A. Filinski. A functional abstraction of typed contexts. Technical Report 89/12, University of Copenhagen, 1989.

[5] Doctor JS: A set of static-analysis tools for JavaScript. https://github.com/mozilla/doctorjs.

[6] C. Earl, M. Might, and D. Van Horn. Pushdown control-flow analysis of higher-order programs. In *Proceedings of the 2010 Workshop on Scheme and Functional Programming*, 2010.

[7] M. Felleisen. The theory and practice of first-class prompts. In *Proceedings of the Fifteenth Annual ACM Symposium on Principles of Programming Languages*, pages 180–190, 1988.

[8] JavaScript generators. https://developer.mozilla.org/en/JavaScript/New_in_JavaScript/1.7.

[9] D. Kranz. *ORBIT: An Optimizing Compiler for Scheme*. PhD thesis, Yale University, 1988.

[10] D. Kranz, R. Kelsey, J. Rees, P. Hudak, J. Philbin, and N. Adams. Orbit: an optimizing compiler for Scheme. In *Proceedings of the 1986 SIGPLAN Symposium on Compiler Construction*, pages 219–233, 1986.

[11] The Lua programming language. http://www.lua.org.

[12] M. Might. *Environment Analysis of Higher-Order Languages*. PhD thesis, Georgia Institute of Technology, 2007.

[13] M. Might and O. Shivers. Analyzing the environment structure of higher-order languages using frame strings. *Theoretical Computer Science*, 375(1–3):137–168, May 2007.

[14] Python generators. http://www.python.org/dev/peps/pep-0255.

[15] T. W. Reps, S. Horwitz, and S. Sagiv. Precise interprocedural dataflow analysis via graph reachability. In *Proceedings of the Twenty-Second Annual ACM Symposium on Principles of Programming Languages*, pages 49–61, 1995.

[16] T. Rompf, I. Maier, and M. Odersky. Implementing first-class polymorphic delimited continuations by a type-directed selective CPS-transform. In *Proceeding of the 14th ACM SIGPLAN International Conference on Functional Programming*, pages 317–328, 2009.

[17] M. Sharir and A. Pnueli. Two approaches to interprocedural dataflow analysis. In S. Muchnick and N. Jones, editors, *Program Flow Analysis: Theory and Applications*. Prentice Hall, 1981.

[18] O. Shivers. Continuations and threads: Expressing machine concurrency directly in advanced languages. In O. Danvy, editor, *Proceedings of the Second ACM SIGPLAN Workshop on Continuations (CW 1997)*, Technical report BRICS-NS-96-13, University of Århus, pages 2:1–15, Paris, France, Jan. 1997.

[19] O. Shivers and M. Might. Continuations and transducer composition. In *Proceedings of the 2006 ACM SIGPLAN Conference on Programming Language Design and Implementation (PLDI 2006)*, pages 295–307, Ottawa, Canada, June 2006.

[20] M. Sperber, R. K. Dybvig, M. Flatt, A. van Straaten, R. Findler, and J. Matthews, editors. *Revised*^{6} *Report on the Algorithmic Language Scheme*. Cambridge University Press, 2010.

[21] G. L. Steele. Rabbit: A Compiler for Scheme. Master's thesis, MIT, 1978.

[22] D. Vardoulakis and O. Shivers. CFA2: A context-free approach to control-flow analysis. In *Proceedings of the 19th European Symposium on Programming (ESOP 2010)*, volume 6012 of *Lecture Notes in Computer Science*, pages 570–589, Paphos, Cyprus, Mar. 2010. Springer.

[23] D. Vardoulakis and O. Shivers. Ordering multiple continuations on the stack. In *Proceedings of the 2011 ACM SIGPLAN Workshop on Partial Evaluation and Program Manipulation (PEPM 2011)*, pages 13–22, Austin, Texas, Jan. 2011.

[24] D. Vardoulakis and O. Shivers. CFA2: A context-free approach to control-flow analysis. *Logical Methods in Computer Science*, 7(2:3): 1–39, May 2011.

[25] M. Wand. Continuation-based multiprocessing. In *Conference Record of the 1980 LISP Conference*, pages 19–28, Stanford, 1980.