

CFA2: a Context-Free Approach to Control-Flow Analysis

Dimitrios Vardoulakis and Olin Shivers

Northeastern University
{dimvar,shivers}@ccs.neu.edu

January 8, 2010

Technical Report NU-CCIS-10-01

Abstract. In a functional language, the dominant control-flow mechanism is function call and return. Most higher-order flow analyses, including k -CFA, do not handle call and return well: they remember only a bounded number of pending calls because they approximate programs with control-flow graphs. Call/return mismatch introduces precision-degrading spurious control-flow paths and increases the analysis time. We describe CFA2, the first flow analysis with precise call/return matching in the presence of higher-order functions and tail calls. We formulate CFA2 as an abstract interpretation of programs in continuation-passing style and describe a sound and complete summarization algorithm for our abstract semantics. A preliminary evaluation shows that CFA2 gives more accurate data-flow information than 0CFA and 1CFA.

1 Introduction

Higher-order functional programs can be analyzed using analyses such as the k -CFA family [1]. These algorithms approximate the valid control-flow paths through the program as the set of all paths through a finite graph of abstract machine states, where each state represents a program point plus some amount of abstracted environment and control context.

In fact, this is not a particularly tight approximation. The set of paths through a finite graph is a regular language. However, the execution traces produced by recursive function calls are strings in a *context-free language*. Approximating this control flow with regular-language techniques permits execution paths that do not properly match calls with returns. This is particularly harmful when analyzing higher-order languages, since flowing functional values down these spurious paths can give rise to further “phantom” control-flow structure, along which functional values can then flow, and so forth, in a destructive spiral that not only degrades precision but drives up the cost of the analysis.

Pushdown models of programs can match an unbounded number of calls and returns, tightening up the set of possible executions to strings in a context-free language. Such models have long been used for first-order languages. The functional approach of Sharir and Pnueli [2] computes transfer-functions for whole

procedures by composing transfer-functions of their basic blocks. Then, at a call-node these functions are used to compute the data-flow value of the corresponding return-node directly. This “summary-based” technique has seen widespread use [3, 4]. Other pushdown models include Recursive State Machines [5] and Pushdown Systems [6].

In this paper, we propose CFA2, a pushdown model of higher-order programs. Our contributions can be summarized as follows:

- CFA2 is a flow analysis with precise call/return matching that can be used in the compilation of both typed and untyped languages. No existing analysis for functional languages enjoys all of these properties. k -CFA and its variants do not provide call/return matching (section 3.1). Rehof and Fähndrich’s analysis [7] supports limited call/return matching and applies to typed languages only (section 7).
- CFA2 uses a stack and a heap for variable binding. Variable references are looked up in one or the other, depending on where they appear in the source code. As it turns out, most references in typical programs are read from the stack, which results in significant precision gains. Also, CFA2 can filter certain bindings off the stack to sharpen precision (section 4). k -CFA with abstract garbage collection [8] cannot infer that it is safe to remove these bindings. Last, the stack makes CFA2 resilient to syntax changes like η -expansion. It is well known that k -CFA is sensitive to such changes [9, 10].
- We formulate CFA2 as an abstract interpretation of programs in continuation-passing style (CPS). The abstract semantics uses a stack of unbounded height. Hence, the abstract state space is infinite, unlike k -CFA. To analyze the state space, we extend the tabulation algorithm of Reps et al. [3]. The resulting algorithm is a search-based variant of summarization that can handle higher-order functions and tail recursion. Currently, CFA2 does not handle first-class-control operators such as `call/cc` (section 5).
- We have implemented 0CFA, 1CFA and CFA2 in the Twobit Scheme compiler [11]. Our experimental results show that CFA2 is more precise than 0CFA and 1CFA. Also, CFA2 usually visits a smaller state space (section 6).

2 Preliminary definitions and notational conventions

We begin with a description of our CPS language and its small-step semantics. For brevity, we develop the theory of CFA2 in the untyped λ -calculus. Primitive data, explicit recursion and side-effects can be easily added using standard techniques [1, ch. 3] [12, ch. 9]. Compilers that use CPS [13, 14] usually partition the terms in a program in two disjoint sets, the user and the continuation set, and treat user terms differently from continuation terms.

We adopt this partitioning for our language (Fig. 1). Variables, lambdas and calls are given 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 apply an additional syntactic constraint: the only continuation variable that can appear free in the body of a user lambda

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

Fig. 1. Partitioned CPS

$(\lambda_l(u k) call)$ is k . This simple constraint forbids first-class control [15]. Intuitively, we get such a program by CPS-converting a direct-style program without `call/cc`. We refer to this variant of CPS as “Restricted CPS” (RCPS).

We assume that all variables in a program have distinct names. Concrete syntax enclosed in $\llbracket \cdot \rrbracket$ 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. Labels can be split into disjoint sets according to the innermost user lambda that contains them. For example, in the following program, which has three user lambdas, these sets are $\{1, 6, 4, 8\}$, $\{2, 9, 5, 10\}$ and $\{3, 7\}$.

$$\begin{aligned}
 & (\lambda_1(u1 k1) _6((\lambda_2(u2 k2) _9((\lambda_5(u5) _10(k2 u1)) u2)) \\
 & \quad (\lambda_3(u3 k3) _7(k3 u3)) \\
 & \quad (\lambda_4(u4) _8(k1 u4))))
 \end{aligned}$$

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., $LV(8) = \{u1, k1, u4\}$ and $LV(5) = \{u2, k2, u5\}$. We use this map to model stack behavior, because all the continuation lambdas that “belong” to a given user lambda λ_l get closed by extending λ_l ’s stack frame (cf. section 4). Notice that, for any ψ , $LV(\psi)$ contains exactly one continuation variable.

We use two notations for tuples, (e_1, \dots, e_n) and $\langle e_1, \dots, e_n \rangle$, 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., $3 :: \langle 1, 3, -47 \rangle$.

The semantics of RCPS appears in Fig. 2. 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* state, 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 .

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

$$\mathcal{A}_{cs}(g, \beta, ve) \triangleq \begin{cases} (g, \beta) & \text{Lam?}(g) \\ ve(g, \beta(g)) & \text{Var?}(g) \end{cases}$$

UEval to *UApply*:

$$\begin{aligned}
& \langle \llbracket (f e q)^l \rrbracket, \beta, ve, t \rangle \rightarrow (proc, d, c, ve, l :: t) \\
& proc = \mathcal{A}_{cs}(f, \beta, ve) \\
& d = \mathcal{A}_{cs}(e, \beta, ve) \\
& c = \mathcal{A}_{cs}(q, \beta, ve)
\end{aligned}$$

CEval to *CAApply*:

$$\begin{aligned}
& \langle \llbracket (q e)^\gamma \rrbracket, \beta, ve, t \rangle \rightarrow (proc, d, ve, \gamma :: t) \\
& proc = \mathcal{A}_{cs}(q, \beta, ve) \\
& d = \mathcal{A}_{cs}(e, \beta, ve)
\end{aligned}$$

UApply to *Eval*:

$$\begin{aligned}
& (proc, d, c, ve, t) \rightarrow (call, \beta', ve', t) \\
& proc = \langle \llbracket (\lambda_l (u k) call) \rrbracket, \beta \rangle \\
& \beta' = \beta[u \mapsto t][k \mapsto t] \\
& ve' = ve[(u, t) \mapsto d][(k, t) \mapsto c]
\end{aligned}$$

CAApply to *Eval*:

$$\begin{aligned}
& (proc, d, ve, t) \rightarrow (call, \beta', ve', t) \\
& proc = \langle \llbracket (\lambda_\gamma (u) call) \rrbracket, \beta \rangle \\
& \beta' = \beta[u \mapsto t] \\
& ve' = ve[(u, t) \mapsto d]
\end{aligned}$$

Fig. 2. Concrete semantics and domains for Restricted CPS

Let's look at the transitions more closely. At a *UEval* state with call site $(f e q)^l$, we evaluate f , e and q using the function \mathcal{A}_{cs} . 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.

From *UApply* to *Eval*, we bind the formals of a procedure $\langle \llbracket (\lambda_l (u k) call) \rrbracket, \beta \rangle$ to the arguments and jump to its body. The new binding environment β' is an extension of 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 argument c .

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 $((pr, \emptyset), input, halt, \emptyset, \langle \rangle)$, where *input* is a closure of the form $\langle \llbracket (\lambda_l (u k) call) \rrbracket, \emptyset \rangle$. Note that the initial time is the empty sequence of call sites.

In the terminology of abstract interpretation, 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. CFA2 and k -CFA are such approximations.

CPS-based compilers may or may not use a control stack for the final code. Steele's view, illustrated in the Rabbit compiler [13], is that argument evaluation pushes stack and function calls are GOTOs. Since arguments in CPS are not calls,

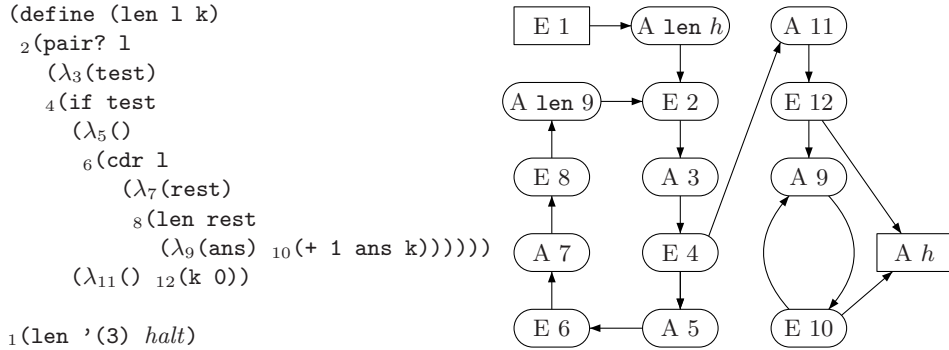


Fig. 3. OCFA on len

argument evaluation is always trivial and Rabbit never needs to push stack. By this approach, every call in CPS is a tail call.

An alternative style was used in the Orbit compiler [14]. At every function call, Orbit pushes a frame for the arguments. By this approach, tail calls are only the calls where the continuation argument is a variable. These CPS call sites were in tail position in the initial direct-style program. *CEval* states where the operator is a variable are calls to the current continuation with a return value. Orbit pops the stack at tail calls and before calling the current continuation.

We will see later that the abstract semantics of CFA2 uses a stack, like Orbit. However, CFA2 computes safe flow information which can be used by both aforementioned approaches. The workings of the abstract interpretation are independent of what style an implementor chooses for the final code.

3 Limitations of k -CFA

In this section, we discuss the main causes of imprecision and inefficiency in k -CFA. Our motivation in developing CFA2 is to create a higher-order flow analysis that overcomes these limitations.

We assume some familiarity with k -CFA, and abstract interpretation in general. Detailed descriptions on these topics can be found in [1, 12]. We use Scheme syntax for our example programs.

3.1 k -CFA does not match calls and returns

In order to make the state space of k -CFA finite, Shivers chose a mechanism similar to the call-strings of Sharir and Pnueli [2]. Thus, recursive programs introduce approximation by folding an unbounded number of recursive calls down to a fixed-size call-string. In effect, by applying k -CFA on a higher-order program, we turn it into a finite-state machine. Taken to the extreme, when k is zero, a function can return to any of its callers, not just to the last one.

For example, consider the function that computes the length of a list, written in CPS (Fig. 3). 0CFA on `len` produces the graph in Fig. 3. *Eval* states (marked with “E”) mention the corresponding call site. *Apply* states are marked with “A”. *UApply* states mention the callee and the continuation argument. The continuation variable `k` is bound to either `halt` or λ_9 . The cycle on the left is taken when the test is true, and it leads to a recursive call. The cycle on the right is taken by returning to λ_9 after a recursive call. Every path from the start to the end node is a valid 0CFA execution. In particular, we cannot exclude the path that recurs four times but applies λ_9 twice. By following such a path, the program will terminate with a *non-empty* stack. It is clear that *k*-CFA cannot help much with optimizations that require accurate calculation of the stack change between program states, such as stack allocation of closure environments.

Spurious flows caused by call/return mismatch affect traditional data-flow information as well. For instance, 0CFA-constant-propagation for the program below cannot spot that `n2` is the constant 2, because 1 also flows to `x` and is mistakenly passed to the continuation λ_2 . 1CFA helps in this example, but repeated η -expansion of the identity function can trick *k*-CFA for any *k*.

```
(let ((id ( $\lambda(x\ k)$  (k x))))
      (id 1 ( $\lambda_1(n1)$  (id 2 ( $\lambda_2(n2)$  (+ n1 n2 halt))))))
```

In a non-recursive program, a large enough *k* can provide accurate call/return matching, but this is not desirable because the analysis becomes intractably slow even when *k* is 1 [10]. Moreover, the ubiquity of recursion in higher-order programs calls for a static analysis that can match an unbounded number of calls and returns. This can be done if we approximate programs using pushdown models instead of finite-state machines.

3.2 The environment problem and fake rebinding

In higher-order languages, many bindings of the same variable can be simultaneously live. Determining at compile time whether two references to some variable will be bound in the same run-time environment is referred to as the *environment problem* [1]. For example, trace through the execution of the following direct-style code:

```
(let ((f ( $\lambda(x\ thunk)$  (if (integer? x) (thunk) ( $\lambda_1()$  x))))
      (f 0 (f "foo" "bar")))
```

In the inner call to `f`, `x` is bound to “foo” and λ_1 is returned. We call `f` again; this time, `x` is an integer, so we jump through `(thunk)` to $(\lambda_1() x)$, and reference `x`, which, despite the just-completed test, is *not* an integer: it is the earlier-bound string “foo”. Thus, during abstract interpretation, it is generally *unsafe* to assume that a variable reference has some property just because an earlier reference had that property.

This has an unfortunate consequence: sometimes an earlier reference provides *safe* information about the reference at hand and *k*-CFA does not spot it:

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

In `compose-same`, both references to `f` are always bound at the same time. However, if multiple closures flow to `f`, k -CFA may call one closure at call site 1 and a different closure at call site 2. This flow never happens at run time.

CFA2 tackles this problem by treating references for a variable v differently from one another, depending on their location in the source code. If v appears in a static context where we know the current stack frame is its environment record, we can be precise. If v appears free in some possibly escaping lambda, we cannot predict its extent so we fall back to a conservative approximation.

3.3 Imprecision increases the running time of the analysis

It is known that k -CFA for $k > 0$ is not a cheap analysis, both in theory [10] and in practice [16]. Counterintuitively, imprecision in higher-order control-flow analyses can increase their running time: imprecision induces spurious control-flow paths, along which the analysis must then flow data, thus creating further spurious paths, and so on, in a vicious cycle which creates extra work whose only function is to degrade the precision of the analysis. This is why techniques that aggressively prune the search space, such as abstract garbage collection [8], not only increase the precision, but can also improve the speed of the analysis.

In the previous subsections, we saw examples of information known at compile time that k -CFA cannot exploit. CFA2 uses this information. The enhanced precision of CFA2 has a positive effect on its running time (cf. section 6).

4 The CFA2 semantics

4.1 Abstract semantics

The CFA2 semantics is an abstract interpreter that executes a program in RCPS, using a stack for variable binding and return-point information.

We describe the stack-management policy with an example. Assume that we run the `len` program of section 3. When calling `(len '(3) halt)` we push a frame `[1 ↦ (3)][k ↦ halt]` on the stack. The test `(pair? 1)` is true, so we add the binding `[test ↦ true]` to the top frame and jump to the true branch. We take the `cdr` of `1` and add the binding `[rest ↦ ()]` to the top frame. We call `len` again, push a new frame for its arguments and jump to its body. This time the test is false, so we extend the top frame with `[test ↦ false]` and jump to the false branch. The call to `k` is a function return, so we pop a frame and pass `0` to λ_9 . Call site 10 is also a function return, so we pop the remaining frame and pass `1` to the top-level continuation `halt`.

In general, we push a frame at function entries and pop at tail calls and at function returns. Results of intermediate computations are stored in the top frame. This policy enforces two invariants about the abstract interpreter. First, when executing inside a user function $(\lambda_l (u k) call)$, the domain of the top frame

\widehat{UEval} to \widehat{UApply} :

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

\widehat{UApply} to \widehat{Eval} :

$$\begin{aligned} & (\llbracket (\lambda_l (u \ k) \ call) \rrbracket, \hat{d}, \hat{c}, st, h) \rightsquigarrow (call, st', h') \\ & st' = push([u \mapsto \hat{d}][k \mapsto \hat{c}], st) \\ & h' = \begin{cases} h \sqcup [u \mapsto \hat{d}] & H_?(u) \\ h & S_?(u) \end{cases} \end{aligned}$$

\widehat{CEval} to $\widehat{CAApply}$:

$$\begin{aligned} & (\llbracket (q \ e)^\gamma \rrbracket, st, h) \rightsquigarrow (clam, \hat{d}, st', h) \\ & clam = \mathcal{A}_k(q, st) \\ & \hat{d} = \mathcal{A}_u(e, st, h) \\ & st' = \begin{cases} pop(st) & Var_?(q) \\ st & Lam_?(q) \end{cases} \end{aligned}$$

$\widehat{CAApply}$ to \widehat{Eval} :

$$\begin{aligned} & (\llbracket (\lambda_\gamma (u) \ call) \rrbracket, \hat{d}, st, h) \rightsquigarrow (call, st', h') \\ & st' = setTop([u \mapsto \hat{d}], st) \\ & h' = \begin{cases} h \sqcup [u \mapsto \hat{d}] & H_?(u) \\ h & S_?(u) \end{cases} \end{aligned}$$

$$\begin{aligned} \mathcal{A}_u(e, st, h) & \triangleq \begin{cases} \{e\} & Lam_?(e) \\ st(e) & S_?(e) \\ h(e) & H_?(e) \end{cases} \\ \mathcal{A}_k(q, st) & \triangleq \begin{cases} q & Lam_?(q) \\ st(q) & Var_?(q) \end{cases} \end{aligned}$$

Abstract domains:

$$\begin{aligned} \hat{c} \in \widehat{UEval} & = UCall \times Stack \times Heap \\ \hat{c} \in \widehat{UApply} & = ULam \times \widehat{UClos} \times \widehat{CClos} \times Stack \times Heap \\ \hat{c} \in \widehat{CEval} & = CCall \times Stack \times Heap \\ \hat{c} \in \widehat{CAApply} & = \widehat{CClos} \times \widehat{UClos} \times Stack \times Heap \\ \hat{d} \in \widehat{UClos} & = Pow(ULam) \\ \hat{c} \in \widehat{CClos} & = CLam + halt \\ fr, tf \in Frame & = (UVar \rightarrow \widehat{UClos}) \cup (CVar \rightarrow \widehat{CClos}) \\ st \in Stack & = Frame^* \\ h \in Heap & = UVar \rightarrow \widehat{UClos} \end{aligned}$$

Stack operations:

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

Fig. 4. Abstract semantics and relevant definitions

is a subset of $LV(l)$. Second, the frame below the top frame is the environment of the current continuation.

Each variable v in our example was looked up in the top frame, because each lookup happened while we were executing inside the lambda that binds v . This is not always the case; in the first snippet of section 3.2 there is a reference to x inside λ_1 . When control reaches that reference, the top frame does not belong

to the user lambda that binds x . CFA2 uses a *heap* to look up such references. The following definition makes these concepts precise.

Definition 1 (Stack and heap references).

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

Put differently, if the innermost user lambda that contains ψ is the one that binds v , then v is a stack reference. In addition, if v is bound by a continuation lambda λ_γ , and the innermost user lambda that contains ψ also contains λ_γ , then v is a stack reference. Intuitively, only heap references may escape. We look up stack references in the top frame, and heap references in the heap. Stack lookups below the top frame never happen.

The CFA2 semantics appears in Fig. 4. An abstract value is either an abstract user closure (member of the set \widehat{UClos}) or an abstract continuation closure (member of \widehat{CClos}). An abstract user closure is a set of user lambdas. An abstract continuation closure 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. It contains only user bindings because in RCPS every continuation variable is a stack variable.

On transition from a \widehat{UEval} state $\hat{\zeta}$ to a \widehat{UApply} state $\hat{\zeta}'$, we first evaluate f , e and q . We evaluate user terms using \mathcal{A}_u and continuation terms using \mathcal{A}_k . We non-deterministically choose one of the lambdas that flow to f as the operator in $\hat{\zeta}'$. 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 continuation 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 on the top frame to be $\{ulam\}$, possibly forgetting other lambdas that may flow to f . This “stack filtering” prevents fake rebinding (cf. section 3.2): when we return to $\hat{\zeta}$, we may reach more stack references of f . These references and the current one are all bound at the same time. Since we are committing to $ulam$ in this transition, these references must also be bound to $ulam$.

In the \widehat{UApply} -to- \widehat{Eval} transition, we push a frame for the procedure’s arguments. In addition, if u is a heap variable we must update its binding in the heap. The join operation \sqcup is defined in the usual way.

In a \widehat{CEval} -to- \widehat{CApply} transition, 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 variable, 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 $clam$. When q is a lambda,

$$\begin{aligned}
|(\llbracket (h_1 \dots h_n)^\psi \rrbracket, \beta, ve, t)|_{ca} &= (\llbracket (h_1 \dots h_n)^\psi \rrbracket, toStack(LV(\psi), \beta, ve), |ve|_{ca}) \\
|(\llbracket (\lambda_l (u k) call) \rrbracket, \beta), d, c, ve, t)|_{ca} &= (\llbracket (\lambda_l (u k) call) \rrbracket, |d|_{ca}, |c|_{ca}, st, |ve|_{ca}) \\
\text{where } st &= \begin{cases} \langle \rangle & c = halt \\ toStack(LV(\gamma), \beta', ve) & c = (\llbracket (\lambda_\gamma (u') call') \rrbracket, \beta') \end{cases} \\
|(\llbracket (\lambda_\gamma (u) call) \rrbracket, \beta), d, ve, t)|_{ca} &= (\llbracket (\lambda_\gamma (u) call) \rrbracket, |d|_{ca}, toStack(LV(\gamma), \beta, ve), |ve|_{ca}) \\
|halt, d, ve, t)|_{ca} &= (halt, |d|_{ca}, \langle \rangle, |ve|_{ca}) \\
|(\llbracket (\lambda_l (u k) call) \rrbracket, \beta)|_{ca} &= \{\llbracket (\lambda_l (u k) call) \rrbracket\} \\
|(\llbracket (\lambda_\gamma (u) call) \rrbracket, \beta)|_{ca} &= \llbracket (\lambda_\gamma (u) call) \rrbracket \\
|halt|_{ca} &= halt \\
|ve|_{ca} &= \{(u, \sqcup_t |ve(u, t)|_{ca}) : H_\gamma(u)\} \\
toStack(\{u_1, \dots, u_n, k\}, \beta, ve) &\triangleq \\
\begin{cases} \langle [u_i \mapsto \hat{d}_i][k \mapsto halt] \rangle & \hat{d}_i = |ve(u_i, \beta(u_i))|_{ca} \wedge \\ & halt = ve(k, \beta(k)) \\ \overline{[u_i \mapsto \hat{d}_i][k \mapsto \llbracket (\lambda_\gamma (u) call) \rrbracket]} :: toStack(LV(\gamma), \beta', ve) & \hat{d}_i = |ve(u_i, \beta(u_i))|_{ca} \wedge \\ & (\llbracket (\lambda_\gamma (u) call) \rrbracket, \beta') = ve(k, \beta(k)) \end{cases}
\end{aligned}$$

Fig. 5. From concrete states to abstract states

it is a newly created closure thus the stack does not change. Note that the transition is deterministic, unlike $\widehat{UEval\text{-}to\text{-}UApply}$. Since we always know which continuation we are about to call, call/return mismatch *never* happens. For instance, the function `len` may be called from many places in a program, so multiple continuations may flow to `k`. But, by retrieving `k`'s value from the stack, we always return to the correct continuation.

In the $\widehat{CAApply\text{-}to\text{-}Eval}$ transition, our stack policy dictates that 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.

4.2 Correctness of CFA2

In this section, we show that the CFA2 semantics safely approximates the concrete semantics. First, we define a map $|\cdot|_{ca}$ from concrete to abstract states. Next, we show that if a state ς transitions to ς' in the concrete semantics, the abstract counterpart $|\varsigma|_{ca}$ of ς transitions to a state ζ' which approximates $|\varsigma'|_{ca}$. By proving this, we ensure that the possible behaviors of the abstract interpreter include the actual run-time behavior of the program.

The map $|\cdot|_{ca}$ appears in Fig. 5. The abstraction of an *Eval* state ς of the form $(\llbracket (h_1 \dots h_n)^\psi \rrbracket, \beta, ve, t)$ 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 ζ' over λ_l . According to our stack policy, the top frame must contain bindings for the formals of λ_l and any temporaries added along the path from ζ' to $\hat{\zeta}$. Therefore, the domain of the top frame is a subset of $LV(l)$, i.e., a subset of $LV(\psi)$. For each user variable $u_i \in (LV(\psi) \cap \text{dom}(\beta))$, 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 $(\llbracket (\lambda_\gamma(u) \text{ call}) \rrbracket, \beta')$, 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 the top-level continuation.

The abstraction of a $UApply$ state over $\langle \llbracket (\lambda_l(u k) \text{ call}) \rrbracket, \beta \rangle$ is a \widehat{UApply} state $\hat{\zeta}$ whose operator is $\llbracket (\lambda_l(u k) \text{ call}) \rrbracket$. The stack of $\hat{\zeta}$ is the stack 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.

We can now state our simulation theorem. The proof proceeds by case analysis on the concrete transition relation. The relation $\hat{\zeta}_1 \sqsubseteq \hat{\zeta}_2$ is a partial ordering on abstract states and can be read as “ $\hat{\zeta}_1$ is more precise than $\hat{\zeta}_2$ ”. The proof and the definition of \sqsubseteq can be found in the appendix.

Theorem 1 (Simulation). *If $\zeta \rightarrow \zeta'$ and $|\zeta|_{ca} \sqsubseteq \hat{\zeta}$, then there exists $\hat{\zeta}'$ such that $\hat{\zeta} \rightsquigarrow \hat{\zeta}'$ and $|\zeta'|_{ca} \sqsubseteq \hat{\zeta}'$.*

5 Computing CFA2

In the previous section we saw how CFA2 addresses the problems of k -CFA, but did not discuss how to explore its state space. Since the size of the stack is unbounded, the state space of CFA2 is infinite and the standard workset algorithms for k -CFA [1, 12] will diverge. For this reason, we have designed a new algorithm based on summarization, a dynamic-programming technique widely used in the interprocedural analysis of first-order programs [2–4] and in context-free language (CFL) reachability algorithms [17].

The difficulty with analyzing programs in a way that respects call/return matching is that the reachable program points from a point n do not depend solely on n , but on the stack contents as well. The intuition behind summarization is to flow facts from n with an *empty* stack to another point n' in the same procedure. We say that n' is *same-context reachable* from n . These facts are then suitably combined to get flow facts for the whole program.

Let’s do the simplest data-flow analysis for the *first-order* program of Fig. 6, namely find which nodes are reachable from the entry of the main function. We will do so by using *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 represent intraprocedural paths, hence the name. We write n_f for the entry node and x_f




Fig. 6. Interprocedural flow-graph for a simple program

for the exit node of a procedure f . Solid arrows are intraprocedural steps. Dotted arrows go from call nodes to the corresponding return nodes. Dashed arrows go from call nodes to entries and from exits to return nodes.

We first scan the program to identify the call sites of each procedure and then start the reachability analysis. Obviously, from 1 we can go to 2 and then to 3, so we record $\langle 1, 1 \rangle$, $\langle 1, 2 \rangle$ and $\langle 1, 3 \rangle$. Then 3 calls `sum`, so we jump to its body. Analysing `sum` produces $\langle 9, 9 \rangle$, $\langle 9, 10 \rangle$ and $\langle 9, 11 \rangle$. Node 11 is an exit reachable from 9, so each caller of `sum` can reach its corresponding return point. We keep track of this fact by recording the *summary* edges $\langle 3, 4 \rangle$ and $\langle 6, 7 \rangle$. Now 4 is reachable from 1, so we discover a new path edge $\langle 1, 4 \rangle$. We go on to discover $\langle 1, 5 \rangle$ and $\langle 1, 6 \rangle$. Reachability inside `sum` does not depend on its calling context, so from the summary edge $\langle 6, 7 \rangle$ we infer that we can reach 7, so we add $\langle 1, 7 \rangle$ to the set of path edges. Finally, we record $\langle 1, 8 \rangle$ which is the end of the program.

We cannot apply summarization to higher-order languages out of the box, because we do not know the call sites of a function by looking at a program's source code. We need a *search-based* variant of summarization, that records callers as it discovers them. Specifically, in the previous example we can record the call $\langle 3, 9 \rangle$ when we reach 3. On reaching 11, we record the summary edge $\langle 9, 11 \rangle$. To find possible return points for `sum`, we look at the set of callers. Since 3 calls 9, 11 can return to 4. Later, when we reach 6, we look at the set of summaries and see that `sum` reaches its exit, so 6 can reach 7. Note that our search-based variant of summarization uses entry-to-exit summaries instead of call-to-return summaries.

5.1 Local semantics

Summarization-based algorithms operate on a finite set of program points. Hence, we cannot use (an infinite number of) abstract states as program points. For this reason, we introduce *local states* and define a map $|\cdot|_{al}$ from abstract to local states (Fig. 7). Intuitively, a local state is like an abstract state but with a single frame instead of a stack. Discarding the rest of the stack makes the local state space finite; keeping the top frame allows precise lookups for stack references.

Essentially, the local semantics describes executions that do not touch the rest of the stack (in other words, executions where functions do not return).

$$\begin{array}{l}
\widetilde{UEval} \text{ to } \widetilde{UApply}: \\
(\llbracket (f \ e \ q)^l \rrbracket, tf, h) \approx (ulam, \hat{d}, h) \\
ulam \in \hat{A}_u(f, tf, h) \\
\hat{d} = \hat{A}_u(e, tf, h) \\
\\
\widetilde{UApply} \text{ to } \widetilde{Eval}: \\
(\llbracket (\lambda_l (u \ k) \ call) \rrbracket, \hat{d}, h) \approx (call, [u \mapsto \hat{d}], h') \\
h' = \begin{cases} h \sqcup [u \mapsto \hat{d}] & H_\gamma(u) \\ h & S_\gamma(u) \end{cases} \\
\\
\widetilde{CEval} \text{ to } \widetilde{CAApply}: \\
(\llbracket (clam \ e)^\gamma \rrbracket, tf, h) \approx (clam, \hat{d}, tf, h) \\
\hat{d} = \hat{A}_u(e, tf, h) \\
\\
\widetilde{CAApply} \text{ to } \widetilde{Eval}: \\
(\llbracket (\lambda_\gamma (u) \ call) \rrbracket, \hat{d}, tf, h) \approx (call, tf', h') \\
tf' = tf[u \mapsto \hat{d}] \\
h' = \begin{cases} h \sqcup [u \mapsto \hat{d}] & H_\gamma(u) \\ h & S_\gamma(u) \end{cases}
\end{array}
\qquad
\hat{A}_u(e, tf, h) \triangleq \begin{cases} \{e\} & Lam_\gamma(e) \\ tf(e) & S_\gamma(e) \\ h(e) & H_\gamma(e) \end{cases}$$

$$\begin{array}{l}
\text{Local domains:} \\
\tilde{\zeta} \in \widetilde{Eval} = Call \times \widetilde{Stack} \times Heap \\
\tilde{\zeta} \in \widetilde{UApply} = ULam \times \widetilde{UClos} \times Heap \\
\tilde{\zeta} \in \widetilde{CAApply} = \widetilde{CClos} \times \widetilde{UClos} \times \widetilde{Stack} \times Heap \\
\\
\widetilde{Frame} = UVar \rightarrow \widetilde{UClos} \\
\widetilde{Stack} = \widetilde{Frame} + \langle \rangle \\
\\
\text{Abstract to local maps:} \\
|call, st, h|_{al} = (call, |st|_{al}, h) \\
|ulam, \hat{d}, \hat{c}, st, h|_{al} = (ulam, \hat{d}, h) \\
|(\hat{c}, \hat{d}, st, h)|_{al} = (\hat{c}, \hat{d}, |st|_{al}, h) \\
|tf :: st'|_{al} = \{(u, tf(u)) : UVar_\gamma(u)\} \\
|\langle \rangle|_{al} = \langle \rangle
\end{array}$$

Fig. 7. Local semantics

Thus, a \widetilde{CEval} state with call site $(ke)^\gamma$ has no successor in this semantics. Since functions do not call their continuations, the frames of local states contain only user bindings. Local steps are otherwise similar to abstract steps. The metavariable $\tilde{\zeta}$ ranges over local states. We define the map $|\cdot|_{cl}$ from concrete to local states to be $|\cdot|_{al} \circ |\cdot|_{ca}$.

We can now see the emerging connection between local semantics and summarization: the local semantics is used for intraprocedural steps and function calls, and we discover return points by recording callers and summary edges.

Next, our algorithm needs to distinguish between different kinds of local states: entries, exits, calls, returns and inner states. CPS lends itself naturally to such a categorization:

- A \widetilde{UApply} state corresponds to an **entry** node—control is about to enter the body of a function.
- A \widetilde{CEval} state where the operator is a variable is an **exit** node—a function is about to pass its result to its context.
- A \widetilde{UEval} state where the continuation argument is a variable is also an **exit**—at tail calls control does not return to the caller.
- A \widetilde{UEval} state where the continuation argument is a lambda is a **call**.
- A \widetilde{CEval} state where the operator is a lambda is an **inner** state.

- A \widetilde{CAppl} state is a **return** if its predecessor is an exit, or an **inner** state if its predecessor is also an inner state. Our algorithm will not need to distinguish between the two kinds of \widetilde{CAppl} s; the difference is just conceptual.

Last, we generalize the notion of summary edges to handle tail recursion. In the following, we rewrite `sum` in a functional style and place it in a context where it gets called three times:

```
(let ((sum (λ(x y k) (+ x y (λ(z) 1(print z k))))))
...2(sum 2 3 (λ3(u1) call3))...
...4(sum 4 5 (λ5(u2) call5))...
...((λ6(n k2) 7(sum n 1 k2)) 9 (λ8(u3) call8))...)
```

The first time (site 2), we record a summary edge from the entry of `sum` to its exit at call site 1, and return to λ_3 . Then, at the second call (site 4) we use the summary edge to find that `sum` will pass its result to λ_5 . The third call is a tail call, so no continuation is born at call site 7. Upon return from `sum`, we must be careful to pass the result to λ_8 . Also, we must restore the environment of the call to λ_6 , *not* the environment of the tail call. We achieve these by recording a “cross-procedure” summary edge from the entry of λ_6 to call site 1. This transitive nature of summaries is essential for tail recursion.

5.2 Summarization

The algorithm for CFA2 is shown in Fig. 8. It is a search-based summarization for higher-order programs with tail calls. Its goal is to compute which local states are reachable from the initial state of a program through paths that respect call/return matching.

An edge $(\tilde{\zeta}_1, \tilde{\zeta}_2)$ is an ordered pair of local states. We call $\tilde{\zeta}_1$ the *source* and $\tilde{\zeta}_2$ the *target* of the edge. 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 (from an entry to a \widetilde{CEval} exit, not necessarily in the same procedure). The target of an edge in *Seen* is reachable from the source and from the initial state (cf. theorem 2). Summaries are also stored in *Summary*.

The workset *W* contains path edges and summaries to be examined. *Final* records \widetilde{CAppl} states that call *halt* with a return value for the whole program. *Callers* contains triples $\langle \tilde{\zeta}_1, \tilde{\zeta}_2, \tilde{\zeta}_3 \rangle$, where $\tilde{\zeta}_1$ is an entry, $\tilde{\zeta}_2$ is a call in the same procedure and $\tilde{\zeta}_3$ is the entry of the callee. *TCallers* contains triples $\langle \tilde{\zeta}_1, \tilde{\zeta}_2, \tilde{\zeta}_3 \rangle$, where $\tilde{\zeta}_1$ is an entry, $\tilde{\zeta}_2$ is a tail call in the same procedure and $\tilde{\zeta}_3$ is the entry of the callee. The initial state $\tilde{\mathcal{I}}(pr)$ is defined as $|\mathcal{I}(pr)|_{cl}$. The helper function $succ(\tilde{\zeta})$ returns the successor(s) of $\tilde{\zeta}$ according to the local semantics.

At every iteration, we remove an edge $(\tilde{\zeta}_1, \tilde{\zeta}_2)$ from *W* and branch depending on $\tilde{\zeta}_2$. If $\tilde{\zeta}_2$ is an entry, a return or an inner state (line 6), then its successor $\tilde{\zeta}_3$ is a state in the same procedure. Since $\tilde{\zeta}_2$ is reachable from $\tilde{\zeta}_1$, $\tilde{\zeta}_3$ is also reachable from $\tilde{\zeta}_1$. If we have not already recorded the edge $(\tilde{\zeta}_1, \tilde{\zeta}_3)$, we do it now (line 25).

```

01  Summary, Callers, TCallers, Final  $\leftarrow \emptyset$ 
02  Seen, W  $\leftarrow \{\tilde{I}(pr), \tilde{I}(pr)\}$ 
03  while  $W \neq \emptyset$ 
04    remove  $(\tilde{\zeta}_1, \tilde{\zeta}_2)$  from  $W$ 
05    switch  $\tilde{\zeta}_2$ 
06      case  $\tilde{\zeta}_2$  of Entry, CApply, Inner-CEval
07        for each  $\tilde{\zeta}_3$  in succ( $\tilde{\zeta}_2$ ) Propagate( $\tilde{\zeta}_1, \tilde{\zeta}_3$ )
08      case  $\tilde{\zeta}_2$  of Call
09        for each  $\tilde{\zeta}_3$  in succ( $\tilde{\zeta}_2$ )
10          Propagate( $\tilde{\zeta}_3, \tilde{\zeta}_3$ )
11          insert  $(\tilde{\zeta}_1, \tilde{\zeta}_2, \tilde{\zeta}_3)$  in Callers
12          for each  $(\tilde{\zeta}_3, \tilde{\zeta}_4)$  in Summary Update( $\tilde{\zeta}_1, \tilde{\zeta}_2, \tilde{\zeta}_3, \tilde{\zeta}_4$ )
13      case  $\tilde{\zeta}_2$  of Exit-CEval
14        if  $\tilde{\zeta}_1 = \tilde{I}(pr)$  then
15          Final( $\tilde{\zeta}_2$ )
16        else
17          insert  $(\tilde{\zeta}_1, \tilde{\zeta}_2)$  in Summary
18          for each  $(\tilde{\zeta}_3, \tilde{\zeta}_4, \tilde{\zeta}_1)$  in Callers Update( $\tilde{\zeta}_3, \tilde{\zeta}_4, \tilde{\zeta}_1, \tilde{\zeta}_2$ )
19          for each  $(\tilde{\zeta}_3, \tilde{\zeta}_4, \tilde{\zeta}_1)$  in TCallers Propagate( $\tilde{\zeta}_3, \tilde{\zeta}_2$ )
20      case  $\tilde{\zeta}_2$  of Exit-TC
21        for each  $\tilde{\zeta}_3$  in succ( $\tilde{\zeta}_2$ )
22          Propagate( $\tilde{\zeta}_3, \tilde{\zeta}_3$ )
23          insert  $(\tilde{\zeta}_1, \tilde{\zeta}_2, \tilde{\zeta}_3)$  in TCallers
24          for each  $(\tilde{\zeta}_3, \tilde{\zeta}_4)$  in Summary Propagate( $\tilde{\zeta}_1, \tilde{\zeta}_4$ )

Propagate( $\tilde{\zeta}_1, \tilde{\zeta}_2$ )  $\triangleq$ 
25  if  $(\tilde{\zeta}_1, \tilde{\zeta}_2)$  not in Seen then insert  $(\tilde{\zeta}_1, \tilde{\zeta}_2)$  in Seen and  $W$ 

Update( $\tilde{\zeta}_1, \tilde{\zeta}_2, \tilde{\zeta}_3, \tilde{\zeta}_4$ )  $\triangleq$ 
26   $\tilde{\zeta}_1$  of the form  $(\llbracket (\lambda_{l_1} (u_1 k_1) call_1) \rrbracket, \hat{d}_1, h_1)$ 
27   $\tilde{\zeta}_2$  of the form  $(\llbracket (f e_2 (\lambda_{\gamma_2} (u_2) call_2))^{l_2} \rrbracket, tf_2, h_2)$ 
28   $\tilde{\zeta}_3$  of the form  $(\llbracket (\lambda_{l_3} (u_3 k_3) call_3) \rrbracket, \hat{d}_3, h_2)$ 
29   $\tilde{\zeta}_4$  of the form  $(\llbracket (k_4 e_4)^{\gamma_4} \rrbracket, tf_4, h_4)$ 
30   $\hat{d} \leftarrow \hat{A}_u(e_4, tf_4, h_4)$ 
31   $tf \leftarrow \begin{cases} tf_2[f \mapsto \{\llbracket (\lambda_{l_3} (u_3 k_3) call_3) \rrbracket\}] & S_{\gamma}(f) \\ tf_2 & H_{\gamma}(f) \vee Lam_{\gamma}(f) \end{cases}$ 
32   $\tilde{\zeta} \leftarrow (\llbracket (\lambda_{\gamma_2} (u_2) call_2) \rrbracket, \hat{d}, tf, h_4)$ 
33  Propagate( $\tilde{\zeta}_1, \tilde{\zeta}$ )

Final( $\tilde{\zeta}$ )  $\triangleq$ 
34   $\tilde{\zeta}$  of the form  $(\llbracket (k e)^{\gamma} \rrbracket, tf, h)$ 
35  insert  $(halt, \hat{A}_u(e, tf, h), \langle \rangle, h)$  in Final

```

Fig. 8. CFA2 workset algorithm

If $\tilde{\zeta}_2$ is a call (line 8) then $\tilde{\zeta}_3$ is an entry of a new procedure, so we propagate $(\tilde{\zeta}_3, \tilde{\zeta}_3)$ instead of $(\tilde{\zeta}_1, \tilde{\zeta}_3)$ (line 10). Next, we record the call in *Callers*. If an exit $\tilde{\zeta}_4$ is reachable from $\tilde{\zeta}_3$, it should return its result to the continuation born at $\tilde{\zeta}_2$ (line 12). 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 29-30). 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 31). The new state $\tilde{\zeta}$ is the corresponding return node of $\tilde{\zeta}_2$, so we propagate $(\tilde{\zeta}_1, \tilde{\zeta})$ (lines 32-33).

If $\tilde{\zeta}_2$ is a \widetilde{CEval} exit and $\tilde{\zeta}_1$ is the initial state (lines 14-15), then $\tilde{\zeta}_2$'s successor is a final state (lines 34-35). If $\tilde{\zeta}_1$ is some other entry, we record the edge in *Summary* and pass the result of $\tilde{\zeta}_2$ to the callers of $\tilde{\zeta}_1$ (lines 17-18). Last, consider the case of a tail call $\tilde{\zeta}_4$ to $\tilde{\zeta}_1$ (line 19). No continuation is born at $\tilde{\zeta}_4$. Thus, we must find where $\tilde{\zeta}_3$ (the entry that led to the tail call) was called from. Then again, it is possible that all calls to $\tilde{\zeta}_3$ are tail calls, in which case we keep searching further back in the call chain to find a return point. We do this backward search by transitively adding a summary edge from $\tilde{\zeta}_3$ to $\tilde{\zeta}_2$ (line 25).

If $\tilde{\zeta}_2$ is a tail call (line 20), we find its successors and record the call in *TCallers* (lines 21-23). If a successor of $\tilde{\zeta}_2$ goes to an exit, we propagate a summary transitively (line 24).

The local state space is finite, so there is a finite number of 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 terminates.

We obviously cannot visit an infinite number of abstract states. To establish the soundness of our flow analysis, we show that if an abstract state $\hat{\zeta}$ is reachable from the initial state, then the algorithm visits $|\hat{\zeta}|_{al}$ (cf. theorem 2). For instance, CFA2 on **1en** (cf. section 3) will tell us that we reach program point 10, *not* that we reach 10 with a stack of size 1, 2, 3 etc.

Soundness guarantees that the CFA2 algorithm does not miss any flows, but it could also compute flows that do not happen in the abstract semantics. For example, a sound but useless algorithm would add all pairs of local states in *Seen*. We establish the completeness of our algorithm by proving that every visited edge has a corresponding abstract flow (cf. theorem 3).

The theorems use two definitions. The first associates a state $\hat{\zeta}$ with its *corresponding entry*, i.e., the entry of the procedure that contains $\hat{\zeta}$. The second finds all entries that reach the corresponding entry of $\hat{\zeta}$ through tail calls. We include the proofs of the theorems in the appendix.

Definition 2. *The Corresponding Entry $CE_p(\hat{\zeta})$ of a state $\hat{\zeta}$ in a path p is:*

- $\hat{\zeta}$, if $\hat{\zeta}$ is an Entry
- $\hat{\zeta}_1$, if $\hat{\zeta}$ is not an Entry, $\hat{\zeta}_2$ is not an Exit-CEval,
 $p \equiv p_1 \rightsquigarrow \hat{\zeta}_1 \rightsquigarrow^* \hat{\zeta}_2 \rightsquigarrow \hat{\zeta} \rightsquigarrow p_2$, and $CE_p(\hat{\zeta}_2) = \hat{\zeta}_1$
- $\hat{\zeta}_1$, if $\hat{\zeta}$ is not an Entry, $p \equiv p_1 \rightsquigarrow \hat{\zeta}_1 \rightsquigarrow^+ \hat{\zeta}_2 \rightsquigarrow \hat{\zeta}_3 \rightsquigarrow^+ \hat{\zeta}_4 \rightsquigarrow \hat{\zeta} \rightsquigarrow p_2$,
 $\hat{\zeta}_2$ is a Call and $\hat{\zeta}_4$ is an Exit-CEval, $CE_p(\hat{\zeta}_2) = \hat{\zeta}_1$, and $\hat{\zeta}_3 \in CE_p^*(\hat{\zeta}_4)$

Definition 3. *For a state $\hat{\zeta}$ and a path p , $CE_p^*(\hat{\zeta})$ is the smallest set such that:*

- $CE_p(\hat{\zeta}) \in CE_p^*(\hat{\zeta})$
- $CE_p^*(\hat{\zeta}_1) \subseteq CE_p^*(\hat{\zeta})$, when $p \equiv p_1 \rightsquigarrow \hat{\zeta}_1 \rightsquigarrow \hat{\zeta}_2 \rightsquigarrow^* \hat{\zeta} \rightsquigarrow p_2$,
 $\hat{\zeta}_1$ is a Tail Call, $\hat{\zeta}_2$ is an Entry, and $\hat{\zeta}_2 = CE_p(\hat{\zeta})$

Theorem 2 (Soundness). *If $p \equiv \hat{\mathcal{I}}(pr) \rightsquigarrow^* \hat{\zeta}$ then, after summarization:*

- if $\hat{\zeta}$ is not a final state then $(|CE_p(\hat{\zeta})|_{al}, |\hat{\zeta}|_{al}) \in Seen$
- if $\hat{\zeta}$ is a final state then $|\hat{\zeta}|_{al} \in Final$
- if $\hat{\zeta}$ is an Exit-CEval and $\hat{\zeta}' \in CE_p^*(\hat{\zeta})$ then $(|\hat{\zeta}'|_{al}, |\hat{\zeta}|_{al}) \in Seen$

	$S_?$	$H_?$	0CFA		1CFA		CFA2	
			visited	constants	visited	constants	visited	constants
<code>len</code>	9	0	81	0	126	0	55	2
<code>rev-iter</code>	17	0	121	0	198	0	82	4
<code>len-Y</code>	15	4	199	0	356	0	131	2
<code>tree-count</code>	33	0	293	2	2856	6	183	10
<code>ins-sort</code>	33	5	509	0	1597	0	600	4
<code>DFS</code>	94	11	1337	8	6890	8	1719	16
<code>flatten</code>	37	0	1520	0	6865	0	478	5
<code>sets</code>	90	3	3915	0	54414	0	4251	4
<code>church-nums</code>	46	23	19130	0	19411	0	22671	0

Fig. 9. Benchmark results

Theorem 3 (Completeness). *After summarization:*

- For each (ζ_1, ζ_2) in *Seen*, there exist $\hat{\zeta}_1, \hat{\zeta}_2$ and p such that
 $p \equiv \hat{\mathcal{I}}(pr) \rightsquigarrow^* \hat{\zeta}_1 \rightsquigarrow^* \hat{\zeta}_2$ and $\zeta_1 = |\hat{\zeta}_1|_{al}$ and $\zeta_2 = |\hat{\zeta}_2|_{al}$ and $\hat{\zeta}_1 \in CE_p^*(\hat{\zeta}_2)$
- For each ζ in *Final*, there exist $\hat{\zeta}$ and p such that
 $p \equiv \hat{\mathcal{I}}(pr) \rightsquigarrow^+ \hat{\zeta}$ and $\zeta = |\hat{\zeta}|_{al}$ and $\hat{\zeta}$ is a final state.

6 Evaluation

We implemented CFA2, 0CFA and 1CFA for the Twobit Scheme compiler [11] and used them to do constant propagation and folding. In this section we report on some initial measurements and comparisons.

We compared the effectiveness of the analyses on a small set of benchmarks (Fig. 9). We measured the number of stack and heap references in each program and the number of constants found by each analysis. We also recorded what goes in the workset in each analysis, i.e., the number of abstract states visited by 0CFA and 1CFA, and the number of path and summary edges visited by CFA2. The running time of an abstract interpretation is proportional to the amount of things inserted in the workset.

We chose programs that exhibit a variety of control-flow patterns. `Len` computes the length of a list recursively. `Rev-iter` reverses a list tail-recursively. `Len-Y` computes the length of a list using the Y-combinator instead of explicit recursion. `Tree-count` counts the nodes in a binary tree. `Ins-sort` sorts a list of numbers using insertion-sort. `DFS` does depth-first search of a graph. `Flatten` turns arbitrarily nested lists into a flat list. `Sets` defines the basic set operations and tests De Morgan’s laws on sets of numbers. `Church-nums` tests distributivity of multiplication over addition for a few Church numerals.

CFA2 finds the most constants, followed by 1CFA. 0CFA is the least precise. CFA2 is also more efficient at exploring its abstract state space. In five out of nine cases, it visits fewer paths than 0CFA does states. The visited set of CFA2 can be up to 3.2 times smaller (`flatten`), and up to 1.3 times larger (`DFS`) than the visited set of 0CFA. 1CFA is less efficient than both 0CFA (9/9 cases) and

CFA2 (8/9 cases). The visited set of 1CFA can be significantly larger than that of CFA2 in some cases (15.6 times in `tree-count`, 14.4 times in `flatten`, 12.8 times in `sets`).

Naturally, the number of stack references in a program is much higher than the number of heap references; most of the time, a variable is referenced only by the lambda that binds it. Thus, CFA2 uses the precise stack lookups more often than the imprecise heap lookups.

7 Related work

We were particularly influenced by Chaudhuri’s recent paper on subcubic algorithms for recursive state machines [4]. His clear and intuitive description of summarization helped us realize that we can use this technique to explore the state space of CFA2.

Reps et al. [3] used summarization to reduce certain data-flow problems for first-order languages to a graph-reachability problem. Our workset algorithm is a variant of their tabulation algorithm, extended for tail recursion and higher-order functions. The reader may have noticed that CFA2 essentially produces a pushdown system. Then, one may wonder why we designed a new algorithm instead of using an existing one like *post** [6, ch. 3]. The reason is that callers cannot be identified syntactically in higher-order languages. Hence, algorithms that analyze higher-order programs must be based on search. The tabulation algorithm can be changed to use search fairly naturally. It is unclear to us how to do that for *post**. In a way, CFA2 creates a pushdown system and analyzes it *at the same time*, much like what *k*-CFA does with control-flow graphs.

Melski and Reps [18] reduced Heintze’s set-constraints [19] to an instance of CFL reachability, which they solve using summarization. Therefore, their solution has the same precision as 0CFA.

CFL reachability has also been used for points-to analysis of imperative higher-order languages. For instance, Sridharan and Bodík’s points-to analysis for Java [20] uses CFL reachability to match writes and reads to object fields. Precise call/return matching is achieved only for programs without recursive methods. Hind’s survey [21] discusses many other variants of points-to analysis.

Debray and Proebsting [22] used ideas from parsing theory to design an interprocedural analysis for first-order programs with tail calls. They describe control-flow with a context-free grammar. Then, the FOLLOW set of a procedure represents its possible return points. Our approach is quite different on the surface, but similar in spirit; we handle tail calls by computing summaries transitively.

Analyses that match an unbounded number of calls and returns have been neglected by the functional language community. The type-based flow analysis of Rehof and Fähndrich [7] is a notable exception. They encode flow information in a type system and then recast the type inference problem to an instance of CFL reachability. The type system uses let-polymorphism. As a result, it provides precise call/return matching for let- and letrec-bound variables but not

for lambda-bound variables. For instance, if we lambda-bind `id` in our earlier example, their type system will not find `n2` to be constant:

```
((λ(id) (let ((n1 (id 1))
              (n2 (id 2)))
  (+ n1 n2)))
(λ(x) x))
```

CFA2 does not distinguish between `let` and `lambda`; in fact, the AST of Twobit contains no lets.

Midtgaard and Jensen [23] created a flow analysis for direct-style higher-order programs that keeps track of “return flow”. They point out that continuations make return-point information explicit in CPS and show how to recover this information in direct-style programs. Their work does not address the issue of unbounded call/return matching.

Might and Shivers [8] proposed GCFA (abstract garbage collection) and μ CFA (abstract counting) to increase the precision of k -CFA. GCFA removes unreachable bindings from the variable environment, and μ CFA counts how many times a variable is bound during the analysis. The two techniques combined reduce the number of spurious flows and give precise environment information. Stack references in CFA2 have a similar effect, because different calls to the same function use different frames. However, we can utilize GCFA and μ CFA to improve precision in the heap.

Recently, Kobayashi [24] proposed a way to statically verify properties of typed higher-order programs using model-checking. He models a program by a higher-order recursion scheme \mathcal{G} , expresses the property of interest in the modal μ -calculus and checks if the infinite tree generated by \mathcal{G} satisfies the property. This technique can do flow analysis, since flow analysis can be encoded as a model-checking problem. The target language of this work is the simply-typed lambda calculus. Programs in a Turing-complete language must be approximated in the simply-typed lambda calculus in order to be model-checked.

8 Conclusions

In this paper we propose CFA2, a pushdown model of higher-order programs, and prove it correct. CFA2 provides precise call/return matching and has a better approach to variable binding than k -CFA. Our evaluation shows that CFA2 gives more accurate data-flow information than 0CFA and 1CFA.

CFA2 is monovariant in the heap. It can be easily extended with call-strings polyvariance, like k -CFA, to produce a family of analyses CFA2.0, CFA2.1 and so on. Then, any instance of CFA2. k would be strictly more precise than the corresponding instance of k -CFA. Another possibility is to add contours in the style of Agesen [25] or Wright and Jagannathan [9]. Note that CFA2 already has most of the above polyvariance “accidentally”, because of the stack lookups.

We believe that pushdown models are a better tool for higher-order flow analysis than control-flow graphs, and are working on providing more empirical

support to this thesis. We plan to use CFA2 for environment analysis and stack-related optimizations. We also plan to add support for `call/cc` in CFA2.

Acknowledgements. We would like to thank Will Clinger and Felix Klock for their help with Twobit, and Manuel Fähndrich and Naoki Kobayashi for clarifications regarding their work. Comments from Bryan Chadwick, Matthias Felleisen, Felix Klock, Aaron Turon and the anonymous referees of ESOP 2010 greatly improved the presentation of the material.

References

1. Shivers, O.: Control-Flow Analysis of Higher-Order Languages. PhD thesis, Carnegie-Mellon University (1991)
2. Sharir, M., Pnueli, A.: Two Approaches to Interprocedural Data Flow Analysis. In: Program Flow Analysis, Theory and Application. Prentice Hall (1981)
3. Reps, T.W., Horwitz, S., Sagiv, S.: Precise Interprocedural Dataflow Analysis via Graph Reachability. In: Principles of Programming Languages. (1995) 49–61
4. Chaudhuri, S.: Subcubic Algorithms for Recursive State Machines. In: Principles of Programming Languages. (2008) 159–169
5. Alur, R., Benedikt, M., Etessami, K., Godefroid, P., Reps, T.W., Yannakakis, M.: Analysis of Recursive State Machines. Transactions on Programming Languages and Systems **27**(4) (2005) 786–818
6. Schwoon, S.: Model-Checking Pushdown Systems. PhD thesis, Technische Universität München (2002)
7. Rehof, J., Fähndrich, M.: Type-Based Flow Analysis: From Polymorphic Subtyping to CFL-Reachability. In: Principles of Programming Languages. (2001) 54–66
8. Might, M., Shivers, O.: Improving Flow Analyses via *FCFA*: Abstract Garbage Collection and Counting. In: International Conference on Functional Programming. (2006) 13–25
9. Wright, A., Jagannathan, S.: Polymorphic Splitting: An Effective Polyvariant Flow Analysis. Trans. on Programming Languages and Systems **20**(1) (1998) 166–207
10. Van Horn, D., Mairson, H.G.: Deciding *k*-CFA is complete for EXPTIME. In: International Conference on Functional Programming. (2008) 275–282
11. Clinger, W.D., Hansen, L.T.: Lambda, the Ultimate Label or a Simple Optimizing Compiler for Scheme. In: LISP and Functional Programming. (1994) 128–139
12. Might, M.: Environment Analysis of Higher-Order Languages. PhD thesis, Georgia Institute of Technology (2007)
13. Steele, G.L.: Rabbit: A Compiler for Scheme. Master’s thesis, MIT (1978)
14. Kranz, D.: ORBIT: An Optimizing Compiler for Scheme. PhD thesis, Yale University (1988)
15. Sabry, A., Felleisen, M.: Reasoning About Programs in Continuation-Passing Style. In: LISP and Functional Programming. (1992) 288–298
16. Shivers, O.: Higher-Order Control-Flow Analysis in Retrospect: Lessons Learned, Lessons Abandoned. In: Best of PLDI. (2004) 257–269
17. Yannakakis, M.: Graph-Theoretic Methods in Database Theory. In: Principles of Database Systems. (1990) 230–242
18. Melski, D., Reps, T.: Interconvertibility of a Class of Set Constraints and Context-Free-Language Reachability. Theoretical Comp. Sci. **248**(1-2) (2000) 29–98

19. Heintze, N.: Set-based program analysis. PhD thesis, Carnegie-Mellon Univ. (1992)
20. Sridharan, M., Bodík, R.: Refinement-based context-sensitive points-to analysis for java. In: Programming Language Design and Implementation. (2006) 387–400
21. Hind, M.: Pointer analysis: haven't we solved this problem yet? In: Program Analysis For Software Tools and Engineering. (2001) 54–61
22. Debray, S.K., Proebsting, T.A.: Interprocedural Control Flow Analysis of First-Order Programs with Tail-Call Optimization. Transactions on Programming Languages and Systems **19**(4) (1997) 568–585
23. Midtgaard, J., Jensen, T.: Control-flow analysis of function calls and returns by abstract interpretation. In: International Conference on Functional Programming. (2009) 287–298
24. Kobayashi, N.: Types and higher-order recursion schemes for verification of higher-order programs. In: Principles of Programming Languages. (2009) 416–428
25. Agesen, O.: The Cartesian Product Algorithm: Simple and Precise Type Inference of Parametric Polymorphism. In: European Conference on Object-Oriented Programming. (1995) 2–26

A Appendix

In section 2, we mentioned that labels in a program can be split into disjoint sets according to the innermost user lambda that contains them. The “label to label” map $LL(\psi)$ returns the labels that are in the same set as ψ . For example, in the following program, $LL(1) = \{1, 6, 4, 8\}$ and $LL(10) = \{2, 9, 5, 10\}$.

$$(\lambda_1(\mathbf{u1} \ \mathbf{k1}) \ \mathbf{6}((\lambda_2(\mathbf{u2} \ \mathbf{k2}) \ \mathbf{9}((\lambda_5(\mathbf{u5}) \ \mathbf{10}(\mathbf{k2} \ \mathbf{u1})) \ \mathbf{u2})) \\ (\lambda_3(\mathbf{u3} \ \mathbf{k3}) \ \mathbf{7}(\mathbf{k3} \ \mathbf{u3})) \\ (\lambda_4(\mathbf{u4}) \ \mathbf{8}(\mathbf{k1} \ \mathbf{u4}))))))$$

Definition 4. For every term g , the map $BV(g)$ returns the variables bound by lambdas which are subterms of g . The map has a simple inductive definition:

$$BV(\llbracket (\lambda_\psi(v_1 \dots v_n) \text{call}) \rrbracket) = \{v_1, \dots, v_n\} \cup BV(\text{call})$$

$$BV(\llbracket (g_1 \dots g_n)^\psi \rrbracket) = BV(g_1) \cup \dots \cup BV(g_n)$$

$$BV(v) = \emptyset \quad \square$$

The following lemma is a simple consequence of alphatisation and lexical scope.

Lemma 1. A concrete state ς has the form (\dots, ve, t) .

- For any closure $(\llbracket (\lambda_\psi(v_1 \dots v_n) \text{call}) \rrbracket, \beta) \in \text{range}(ve)$, it holds that $\text{dom}(\beta) \cap BV(\llbracket (\lambda_\psi(v_1 \dots v_n) \text{call}) \rrbracket) = \emptyset$.
- If ς is an Eval state with call site $\llbracket (g_1 \dots g_n)^\psi \rrbracket$ and environment β , then $\text{dom}(\beta) \cap BV(\llbracket (g_1 \dots g_n)^\psi \rrbracket) = \emptyset$.
- If ς is an Apply state, for any closure $(\llbracket (\lambda_\psi(v_1 \dots v_n) \text{call}) \rrbracket, \beta)$ in operator or argument position, then $\text{dom}(\beta) \cap BV(\llbracket (\lambda_\psi(v_1 \dots v_n) \text{call}) \rrbracket) = \emptyset$.

Proof It suffices to show that the lemma holds for the initial state and it’s maintained by transition. \square

Theorem 4 (Simulation). If $\varsigma \rightarrow \varsigma'$ and $|\varsigma|_{ca} \sqsubseteq \hat{\varsigma}$, then there exists ζ' such that $\hat{\varsigma} \rightsquigarrow \zeta'$ and $|\zeta'|_{ca} \sqsubseteq \zeta'$.

Proof By cases on the concrete transition.

- a) *UEval to UApply*
- $$(\llbracket (f \ e \ q)^l \rrbracket, \beta, ve, t) \rightarrow (\text{proc}, d, c, ve, l :: t)$$
- $$\text{proc} = \mathcal{A}_{cs}(f, \beta, ve)$$
- $$d = \mathcal{A}_{cs}(e, \beta, ve)$$
- $$c = \mathcal{A}_{cs}(q, \beta, ve)$$

Since $|\varsigma|_{ca} \sqsubseteq \hat{\varsigma}$, $\hat{\varsigma}$ is of the form $(\llbracket (f \ e \ q)^l \rrbracket, st, h)$, where $|ve|_{ca} \sqsubseteq h$, $ts = \text{toStack}(LV(l), \beta, ve)$, $ts \sqsubseteq st$.

The abstract transition is

$$(\llbracket (f \ e \ q)^l \rrbracket, st, h) \rightsquigarrow (f', \hat{d}, \hat{c}, st', h)$$

$$f' \in \mathcal{A}_u(f, st, h)$$

$$\hat{d} = \mathcal{A}_u(e, st, h)$$

$$\begin{aligned}
(call_1, st_1, h_1) \sqsubseteq (call_2, st_2, h_2) & \quad \text{iff} \quad call_1 = call_2 \wedge st_1 \sqsubseteq st_2 \wedge h_1 \sqsubseteq h_2 \\
(ulam_1, \hat{d}_1, \hat{c}_1, st_1, h_1) \sqsubseteq (ulam_2, \hat{d}_2, \hat{c}_2, st_2, h_2) & \quad \text{iff} \quad ulam_1 = ulam_2 \wedge \hat{d}_1 \sqsubseteq \hat{d}_2 \wedge \hat{c}_1 = \hat{c}_2 \wedge st_1 \sqsubseteq st_2 \wedge h_1 \sqsubseteq h_2 \\
(\hat{c}, \hat{d}_1, st_1, h_1) \sqsubseteq (\hat{c}', \hat{d}_2, st_2, h_2) & \quad \text{iff} \quad \hat{c} = \hat{c}' \wedge \hat{d}_1 \sqsubseteq \hat{d}_2 \wedge st_1 \sqsubseteq st_2 \wedge h_1 \sqsubseteq h_2 \\
h_1 \sqsubseteq h_2 & \quad \text{iff} \quad h_1(u) \sqsubseteq h_2(u) \quad \text{for each } u \in \text{dom}(h_1) \\
tf_1 :: st_1 \sqsubseteq tf_2 :: st_2 & \quad \text{iff} \quad tf_1 \sqsubseteq tf_2 \wedge st_1 \sqsubseteq st_2 \\
\langle \rangle \sqsubseteq \langle \rangle & \\
tf_1 \sqsubseteq tf_2 & \quad \text{iff} \quad tf_1(v) \sqsubseteq tf_2(v) \quad \text{for each } v \in \text{dom}(tf_1) \\
\hat{d}_1 \sqsubseteq \hat{d}_2 & \quad \text{iff} \quad \hat{d}_1 \subseteq \hat{d}_2 \\
\hat{c} \sqsubseteq \hat{c} &
\end{aligned}$$

Fig. 10. The \sqsubseteq relation on abstract states

$$\begin{aligned}
\hat{c} &= \mathcal{A}_k(q, st) \\
st' &= \begin{cases} pop(st) & Var_?(q) \\ st & Lam_?(q) \wedge (H_?(f) \vee Lam_?(f)) \\ setTop([f \mapsto \{f'\}], st) & Lam_?(q) \wedge S_?(f) \end{cases}
\end{aligned}$$

We must show that there is a state ζ' such that $|\zeta'|_{ca} \sqsubseteq \zeta'$. The variable environment and the heap remain unchanged in the transitions, so for ζ' and \hat{c}' we know that $|ve|_{ca} \sqsubseteq h$. We proceed by cases for the rest components of ζ .

- $Var_?(q)$ and $c = ve(q, \beta(q)) = halt$
Then, $ts(q) = halt$.
Also, $\hat{c} = st(q) \sqsupseteq ts(q) = halt \Rightarrow \hat{c} = halt$ and $|c|_{ca} \sqsubseteq \hat{c}$.
The stack of $|\zeta'|_{ca}$ is $ts' = \langle \rangle$.
But ts and st have the same size, and $st' = pop(st)$, so $st' = \langle \rangle$, thus $ts' \sqsubseteq st'$.
- $Lam_?(q)$ and $c = ve(q, \beta(q)) = \llbracket (\lambda_\gamma(u_1) call_1) \rrbracket, \beta'$
Then, $ts(q) = \llbracket (\lambda_\gamma(u_1) call_1) \rrbracket$.
Also, $\hat{c} = st(q) \sqsupseteq ts(q) = \llbracket (\lambda_\gamma(u_1) call_1) \rrbracket \Rightarrow \hat{c} = \llbracket (\lambda_\gamma(u_1) call_1) \rrbracket$ and $|c|_{ca} \sqsubseteq \hat{c}$.
The stack of $|\zeta'|_{ca}$ is $ts' = toStack(LV(\gamma), \beta', ve) = pop(ts)$.
But $st' = pop(st)$, and $ts \sqsubseteq st$ which implies $pop(ts) \sqsubseteq pop(st)$, thus $ts' \sqsubseteq st'$.
- $setTop([f \mapsto \{f'\}], st)$
Then, $c = (q, \beta), \hat{c} = q$ and $|c|_{ca} \sqsubseteq \hat{c}$. There are two cases for st' depending on f , and we examine them below.

- $Lam_{\gamma}(f)$
Then, $proc = (f, \beta)$ and $f' = f$, thus $|proc|_{ca} \sqsubseteq \{f'\}$.
Moreover, if $Lam_{\gamma}(q)$ then $ts' = ts$ and $st' = st$, so $ts' \sqsubseteq st'$.
- $H_{\gamma}(f)$
Then, $proc = ve(f, \beta(f))$,
a closure of the form $(\llbracket (\lambda_{l'}(u_2 k_2) call_2) \rrbracket, \beta_2)$.
Therefore, $|proc|_{ca} = \{\llbracket (\lambda_{l'}(u_2 k_2) call_2) \rrbracket\}$.
Since $|ve|_{ca} \sqsubseteq h$ and $\llbracket (\lambda_{l'}(u_2 k_2) call_2) \rrbracket \in |ve|_{ca}(f)$,
we get $\llbracket (\lambda_{l'}(u_2 k_2) call_2) \rrbracket \in h(f)$.
So, we pick f' to be $\llbracket (\lambda_{l'}(u_2 k_2) call_2) \rrbracket$.
Also, if $Lam_{\gamma}(q)$, the argument for the stacks is analogous to the previous case.
- $S_{\gamma}(f)$
Then, $proc = ve(f, \beta(f))$,
a closure of the form $(\llbracket (\lambda_{l'}(u_2 k_2) call_2) \rrbracket, \beta_2)$.
Therefore, $|proc|_{ca} = \{\llbracket (\lambda_{l'}(u_2 k_2) call_2) \rrbracket\}$
and also $ts(f) = \{\llbracket (\lambda_{l'}(u_2 k_2) call_2) \rrbracket\}$,
which implies that $\llbracket (\lambda_{l'}(u_2 k_2) call_2) \rrbracket \in st(f)$.
So, we pick f' to be $\llbracket (\lambda_{l'}(u_2 k_2) call_2) \rrbracket$.
In the case when $Lam_{\gamma}(q)$, we want to show that $ts' \sqsubseteq st'$.
If $q = \llbracket (\lambda_{\gamma}(u_1) call_1) \rrbracket$, then $LV(\gamma) = LV(l)$ and $ts' = ts$.
Also, $st' = setTop([f \mapsto \{\llbracket (\lambda_{l'}(u_2 k_2) call_2) \rrbracket\}], st)$.
Since, $pop(ts) \sqsubseteq pop(st)$ we only need to show that the top frames of ts'
and st' are in \sqsubseteq .
For this it suffices to show that $ts'(f) \sqsubseteq st'(f)$ which holds.
- $Lam_{\gamma}(e) \vee H_{\gamma}(e) \vee S_{\gamma}(e)$
Showing that $|d|_{ca} \sqsubseteq \hat{d}$ is similar to showing it for f' or c .

b) *UApply to Eval*

$$\begin{aligned}
&(proc, d, c, ve, t) \rightarrow (call, \beta', ve', t) \\
&proc = \langle \llbracket (\lambda_l(u k) call) \rrbracket, \beta \rangle \\
&\beta' = \beta[u \mapsto t][k \mapsto t] \\
&ve' = ve[(u, t) \mapsto d][(k, t) \mapsto c]
\end{aligned}$$

Since $|c|_{ca} \sqsubseteq \hat{c}$, \hat{c} is of the form $(\llbracket (\lambda_l(u k) call) \rrbracket, \hat{d}, \hat{c}, st, h)$,
where $|d|_{ca} \sqsubseteq \hat{d}$, $|c|_{ca} = \hat{c}$, $ts \sqsubseteq st$, $|ve|_{ca} \sqsubseteq h$, and

$$ts = \begin{cases} \langle \rangle & c = halt \\ toStack(LV(\gamma), \beta_1, ve) & c = (\llbracket (\lambda_{\gamma}(u_1) call_1) \rrbracket, \beta_1) \end{cases}$$

The abstract transition is

$$\begin{aligned}
&(\llbracket (\lambda_l(u k) call) \rrbracket, \hat{d}, \hat{c}, st, h) \rightsquigarrow (call, st', h') \\
&st' = push([u \mapsto \hat{d}][k \mapsto \hat{c}], st) \\
&h' = \begin{cases} h \sqcup [u \mapsto \hat{d}] & H_{\gamma}(u) \\ h & S_{\gamma}(u) \end{cases}
\end{aligned}$$

We must show that $|c'|_{ca} \sqsubseteq \hat{c}'$ i.e. that $ts' \sqsubseteq st'$ and $|ve'|_{ca} \sqsubseteq h'$, where

ts' is the stack of $|\varsigma'|_{ca}$. The innermost user lambda that contains $call$ is λ_l , therefore $ts' = toStack(LV(l), \beta', ve')$.

We assume that $c = (\llbracket (\lambda_\gamma(u_1) call_1) \rrbracket, \beta_1)$ and that $H_\gamma(u)$ holds, the other cases are simpler. In this case, $|ve'|_{ca}$ is the same as $|ve|_{ca}$ except that $|ve'|_{ca}(u) = |ve|_{ca}(u) \sqcup |d|_{ca}$. Also, $h'(u) = h(u) \sqcup \hat{d}$, thus $|ve'|_{ca} \sqsubseteq h'$.

We know that $LV(l) \sqsubseteq BV(call)$.

We also know that β' contains bindings for u and k , and by lemma 1 it doesn't bind any variables in $BV(call)$.

Thus, the top frame of ts' is $[u \mapsto |d|_{ca}][k \mapsto |c|_{ca}]$.

The top frame of st' is $[u \mapsto \hat{d}][k \mapsto \hat{c}]$, therefore the frames are in \sqsubseteq . Also, $pop(ts') = toStack(LV(\gamma), \beta_1, ve')$ and $pop(st') = st$. Since $ts \sqsubseteq st$, it suffices to show that $pop(ts') = ts$. By the temporal consistency of states (cf. [12] definition 4.4.5), ts' won't contain bindings born at time t because they are younger than all bindings in β_1 . This implies that $pop(ts') = ts$.

- c) *CEval to CApply*
 $(\llbracket (qe)^\gamma \rrbracket, \beta, ve, t) \rightarrow (proc, d, ve, \gamma :: t)$
 $proc = \mathcal{A}_{cs}(q, \beta, ve)$
 $d = \mathcal{A}_{cs}(e, \beta, ve)$

Since $|\varsigma|_{ca} \sqsubseteq \hat{\varsigma}$, $\hat{\varsigma}$ is of the form $(\llbracket (qe)^\gamma \rrbracket, st, h)$, where $|ve|_{ca} \sqsubseteq h$, $ts = toStack(LV(\gamma), \beta, ve)$, $ts \sqsubseteq st$.

The abstract transition is
 $(\llbracket (qe)^\gamma \rrbracket, st, h) \rightsquigarrow (q', \hat{d}, st', h)$
 $q' = \mathcal{A}_k(q, st)$
 $\hat{d} = \mathcal{A}_u(e, st, h)$
 $st' = \begin{cases} pop(st) & Var_?(q) \\ st & Lam_?(q) \end{cases}$

We must show that $|\varsigma'|_{ca} \sqsubseteq \hat{\varsigma}'$ i.e. that $|proc|_{ca} = q'$, $|d|_{ca} \sqsubseteq \hat{d}$, and $ts' \sqsubseteq st'$, where ts' is the stack of $|\varsigma'|_{ca}$. We proceed by cases for the components of ς .

- $Lam_?(q)$ i.e. $q = \llbracket (\lambda_{\gamma'}(u_1) call_1) \rrbracket$
Then, $proc = (q, \beta)$, $q' = q$ and $st' = st$. Thus, $|proc|_{ca} = q'$. Also, $ts' = toStack(LV(\gamma'), \beta, ve)$ and $LV(\gamma') = LV(\gamma) \Rightarrow ts' = ts \Rightarrow ts' \sqsubseteq st'$.
- $Var_?(q)$ and $proc = ve(q, \beta(q)) = (\llbracket (\lambda_{\gamma'}(u_1) call_1) \rrbracket, \beta_1)$
Since $q \in LV(\gamma)$ we get $ts(q) = \llbracket (\lambda_{\gamma'}(u_1) call_1) \rrbracket$. From the latter and $ts \sqsubseteq st \Rightarrow st(q) = \llbracket (\lambda_{\gamma'}(u_1) call_1) \rrbracket \Rightarrow q' = \llbracket (\lambda_{\gamma'}(u_1) call_1) \rrbracket \Rightarrow |proc|_{ca} = q'$.
Also, $ts' = toStack(LV(\gamma'), \beta_1, ve) = pop(ts)$ and $st' = pop(st)$ and $ts \sqsubseteq st \Rightarrow pop(ts) \sqsubseteq pop(st) \Rightarrow ts' \sqsubseteq st'$.
- $Var_?(q)$ and $proc = ve(q, \beta(q)) = halt$
Similar to the previous case.
- $Var_?(e)$ or $Lam_?(e)$
In each of these cases it is simple to show $|d|_{ca} \sqsubseteq \hat{d}$.

d) *CApply to Eval*

This case requires arguments similar to the previous cases. \square

Lemma 2. *On an \widehat{Eval} -to- \widehat{Apply} transition, the stack below the top frame is irrelevant. Formally,*

- If $(\llbracket (f e (\lambda_\gamma (u) call))^l \rrbracket, tf :: st, h) \rightsquigarrow (ulam, \hat{d}, \hat{c}, tf' :: st, h)$
then for any st' ,
 $(\llbracket (f e (\lambda_\gamma (u) call))^l \rrbracket, tf :: st', h) \rightsquigarrow (ulam, \hat{d}, \hat{c}, tf' :: st', h)$
- If $(\llbracket (f e k)^l \rrbracket, tf :: st, h) \rightsquigarrow (ulam, \hat{d}, \hat{c}, st, h)$
then for any st' ,
 $(\llbracket (f e k)^l \rrbracket, tf :: st', h) \rightsquigarrow (ulam, \hat{d}, \hat{c}, st', h)$
- Similarly for the \widehat{CEval} -to- \widehat{CApply} transition. \square

Lemma 3. *On an \widehat{Apply} -to- \widehat{Eval} transition, the stack is irrelevant. Formally,*

- If $(\llbracket (\lambda_l (u k) call) \rrbracket, \hat{d}, \hat{c}, st, h) \rightsquigarrow (call, [u \mapsto \hat{d}][k \mapsto \hat{c}] :: st, h')$
then for any st' ,
 $(\llbracket (\lambda_l (u k) call) \rrbracket, \hat{d}, \hat{c}, st', h) \rightsquigarrow (call, [u \mapsto \hat{d}][k \mapsto \hat{c}] :: st', h')$
- Similarly for the \widehat{CApply} -to- \widehat{Eval} transition, where st' is any non-empty stack. \square

Definition 5 (Push Monotonicity).

Let $p \equiv \hat{\zeta}_e \rightsquigarrow^* \hat{\zeta}$ where $\hat{\zeta}_e$ is an entry with stack st_e . The path p is **push monotonic** iff every transition $\hat{\zeta}_1 \rightsquigarrow \hat{\zeta}_2$ satisfies the property:

If the stack of $\hat{\zeta}_1$ is st_e then the transition can only push the stack, it cannot pop or setTop. \square

The following properties are simple consequences of push monotonicity.

Property 1. The stack of the first state in a push-monotonic path is a suffix of the stack of every other state in the path.

Property 2. In a push-monotonic path, the number of pushes is greater than or equal to the number of pops.

The following lemma associates entries with “same-level reachable” states. A state $\hat{\zeta}$ is same-level reachable from an entry $\hat{\zeta}_e$ if it is in the procedure whose entry is $\hat{\zeta}_e$ or if it is in some procedure that can be reached from $\hat{\zeta}_e$ through tail calls i.e. without growing the stack.

Lemma 4 (Same-level reachability).

Let $\hat{\zeta}_e = (\llbracket (\lambda_l (u k) call) \rrbracket, \hat{d}, \hat{c}, st_e, h_e)$, $\hat{\zeta} = (\dots, st, h)$, and $p \equiv \hat{\zeta}_e \rightsquigarrow^* \hat{\zeta}$ where $\hat{\zeta}_e \in CE_p^*(\hat{\zeta})$. Then,

1. If $\hat{\zeta}$ is an entry, $st = st_e$.

2. If $\hat{\zeta}$ is not an entry,
 - (a) st is of the form $tf :: st_e$, for some frame tf .
 - (b) there exists k' such that $tf(k') = \hat{c}$.
 - (c) if $\hat{\zeta}_e = CE_p(\hat{\zeta})$ then $\text{dom}(tf) \subseteq LV(l)$, $tf(u) \sqsubseteq \hat{d}$ and $tf(k) = \hat{c}$.
 Moreover, if $\hat{\zeta}$ is an \widehat{Eval} over call site ψ then $\psi \in LL(l)$, and if $\hat{\zeta}$ is a \widehat{CApply} over $(\lambda_\gamma(u')\text{call}')$ then $\gamma \in LL(l)$.
3. p is push monotonic.

Proof By induction on the length $|p|$ of p . Note that (3) follows from the form of the stack in (1) and (2), so we won't prove it separately.

- If $|p| = 0$, then $\hat{\zeta} = \hat{\zeta}_e$ so $st = st_e$.
- If $|p| > 0$, there are two cases; either $\hat{\zeta}_e = CE_p(\hat{\zeta})$ or $\hat{\zeta}_e \neq CE_p(\hat{\zeta})$.

a) $\hat{\zeta}_e = CE_p(\hat{\zeta})$

Since $|p| > 0$, $\hat{\zeta}$ is not an entry, so the second or the third branch of the definition of CE_p determine the shape of p .

a1) $p \equiv \hat{\zeta}_e \rightsquigarrow^* \hat{\zeta}' \rightsquigarrow \hat{\zeta}$

Here, the predecessor $\hat{\zeta}'$ of $\hat{\zeta}$ is not a \widehat{CEval} exit, and $\hat{\zeta}_e = CE_p(\hat{\zeta}')$. We proceed by cases on $\hat{\zeta}'$. Note that $\hat{\zeta}'$ cannot be a \widehat{UEval} because then $\hat{\zeta}$ is an entry, so $\hat{\zeta} = CE_p(\hat{\zeta})$, and our assumption that $\hat{\zeta}_e = CE_p(\hat{\zeta})$ breaks.

a1.1) $\hat{\zeta}'$ is an inner \widehat{CEval}

Then, $\hat{\zeta}' = (\llbracket (\lambda_\gamma(u')\text{call}') e' \rrbracket^{\gamma'}, st', h')$. By *IH*, $st' = tf' :: st_e$, $\text{dom}(tf') \subseteq LV(l)$, $tf'(u) \sqsubseteq \hat{d}$, $tf'(k) = \hat{c}$ and $\gamma' \in LL(l)$.

By the abstract semantics,

$\hat{\zeta} = (\llbracket (\lambda_\gamma(u')\text{call}') \rrbracket, \hat{d}', st', h')$ where $\hat{d}' = \mathcal{A}_u(e', st', h')$. We know that $\gamma \in LL(l)$ because $\gamma' \in LL(l)$. Also, the stack is unchanged in the transition. Thus, (2a), (2b) and (2c) hold for $\hat{\zeta}$.

a1.2) $\hat{\zeta}'$ is a \widehat{CApply}

Then, $\hat{\zeta}' = (\llbracket (\lambda_\gamma(u')\text{call}') \rrbracket, \hat{d}', st', h')$. By *IH*, $st' = tf' :: st_e$, $\text{dom}(tf') \subseteq LV(l)$, $tf'(u) \sqsubseteq \hat{d}$, $tf'(k) = \hat{c}$ and $\gamma \in LL(l)$.

By the abstract semantics, $\hat{\zeta} = (\text{call}', st, h)$ where

$st = \text{setTop}([u' \mapsto \hat{d}'], st')$.

So, $st = tf :: st_e$ which satisfies (2a). Also, $tf = tf'[u' \mapsto \hat{d}']$ where $u' \in LV(l)$ because $\gamma \in LL(l)$, and $u' \neq u$ because the program is α -tised. Thus, $\text{dom}(tf) = \text{dom}(tf') \cup \{u'\} \subseteq LV(l)$, and $tf(u) = tf'(u) \sqsubseteq \hat{d}$, and $tf(k) = tf'(k) = \hat{c}$. Last, the label of call' is in $LL(l)$ because $\gamma \in LL(l)$.

a1.3) $\hat{\zeta}'$ is a \widehat{UApply}

Then, $\hat{\zeta}' = \hat{\zeta}_e$ because $\hat{\zeta}_e = CE_p(\hat{\zeta}')$. This case is simple.

a2) $p \equiv \hat{\zeta}_e \rightsquigarrow^+ \hat{\zeta}_2 \rightsquigarrow \hat{\zeta}_3 \rightsquigarrow^+ \hat{\zeta}' \rightsquigarrow \hat{\zeta}$

Here, the third branch of the definition of CE_p determines the shape of p , so $\hat{\zeta}_2$ is a call, $\hat{\zeta}_e = CE_p(\hat{\zeta}_2)$, $\hat{\zeta}'$ is a \widehat{CEval} exit and $\hat{\zeta}_3 \in CE_p^*(\hat{\zeta}')$. By *IH* for $\hat{\zeta}_e \rightsquigarrow^+ \hat{\zeta}_2$ we get

$\hat{\zeta}_2 = (\llbracket (f_2 \ e_2 \ (\lambda_{\gamma_2} (u_2) \ call_2))^{l_2} \rrbracket, st_2, h_2)$, where $st_2 \equiv tf_2 :: st_e$,
 $\text{dom}(tf_2) \subseteq LV(l)$, $tf_2(u) \sqsubseteq \hat{d}$, $tf_2(k) = \hat{c}$ and $l_2 \in LL(l)$.

By the abstract semantics for $\hat{\zeta}_2 \rightsquigarrow \hat{\zeta}_3$ we get

$\hat{\zeta}_3 = (\llbracket (\lambda_{l_3} (u_3 \ k_3) \ call_3) \rrbracket, \hat{d}_3, \hat{c}_3, st_3, h_2)$,

where $\llbracket (\lambda_{l_3} (u_3 \ k_3) \ call_3) \rrbracket \in \mathcal{A}_u(f_2, st_2, h_2)$,

$\hat{d}_3 = \mathcal{A}_u(e_2, st_2, h_2)$, $\hat{c}_3 = \llbracket (\lambda_{\gamma_2} (u_2) \ call_2) \rrbracket$ and

either $st_3 = st_2$, if $(Lam_{\gamma}(f_2) \vee H_{\gamma}(f_2))$ holds,

or $st_3 = \text{setTop}([f_2 \mapsto \{\llbracket (\lambda_{l_3} (u_3 \ k_3) \ call_3) \rrbracket\}], st_2)$, if $S_{\gamma}(f_2)$ holds.

a2.1) $S_{\gamma}(f_2)$

Then, $st_3 = tf_2[f_2 \mapsto \{\llbracket (\lambda_{l_3} (u_3 \ k_3) \ call_3) \rrbracket\}] :: st_e$.

By *IH* for $\hat{\zeta}_3 \rightsquigarrow^+ \hat{\zeta}'$ we get $\hat{\zeta}' = (\llbracket (k' \ e')^{\gamma} \rrbracket, st', h')$,

where $st' = tf' :: st_3$

and $tf'(k') = \llbracket (\lambda_{\gamma_2} (u_2) \ call_2) \rrbracket$.

Thus, by the abstract semantics for $\hat{\zeta}' \rightsquigarrow \hat{\zeta}$ we get

$\hat{\zeta} = (\llbracket (\lambda_{\gamma_2} (u_2) \ call_2) \rrbracket, \hat{d}', st_3, h')$.

Now, $\gamma_2 \in LL(l)$ follows from $l_2 \in LL(l)$.

Also, $st = tf :: st_e$

where $tf = tf_2[f_2 \mapsto \{\llbracket (\lambda_{l_3} (u_3 \ k_3) \ call_3) \rrbracket\}]$.

Then, $\text{dom}(tf) = \text{dom}(tf_2) \cup \{f_2\} \subseteq LV(l)$ because $S_{\gamma}(f_2)$ implies $f_2 \in LV(l)$.

Also, $tf(k) = tf_2(k) = \hat{c}$. Last, we take cases depending on whether u and f_2 are the same variable or not.

• $u = f_2$

$tf(u) = \{\llbracket (\lambda_{l_3} (u_3 \ k_3) \ call_3) \rrbracket\}$

$\subseteq \mathcal{A}_u(f_2, st_2, h_2)$

$= st_2(f_2)$

$= tf_2(f_2)$

$= tf_2(u)$

$\sqsubseteq \hat{d}$

• $u \neq f_2$

$tf(u) = tf_2(u) \sqsubseteq \hat{d}$

a2.2) $Lam_{\gamma}(f_2) \vee H_{\gamma}(f_2)$

This case is simpler than the previous case because $st_3 = st_2$.

b) $\hat{\zeta}_e \neq CE_p(\hat{\zeta})$ (but $\hat{\zeta}_e \in CE_p^*(\hat{\zeta})$)

Then, the second branch of the definition of CE_p^* determines the shape of p ; $p \equiv \hat{\zeta}_e \rightsquigarrow^+ \hat{\zeta}_1 \rightsquigarrow \hat{\zeta}_2 \rightsquigarrow^* \hat{\zeta}$, where $\hat{\zeta}_1$ is a tail call, $\hat{\zeta}_2 = CE_p(\hat{\zeta})$ and $\hat{\zeta}_e \in CE_p^*(\hat{\zeta}_1)$.

By *IH* for $\hat{\zeta}_e \rightsquigarrow^+ \hat{\zeta}_1$ we get $\hat{\zeta}_1 = (\llbracket (f_1 \ e_1 \ k_1)^{l_1} \rrbracket, st_1, h_1)$,

where $st_1 = tf_1 :: st_e$, $tf_1(k_1) = \hat{c}$.

By the abstract semantics,

$\hat{\zeta}_2 = (\llbracket (\lambda_{l_2} (u_2 \ k_2) \ call_2) \rrbracket, \hat{d}_2, \hat{c}, st_e, h_1)$.

b.1) $\hat{\zeta}$ is an entry

Then, $\hat{\zeta} = \hat{\zeta}_2$ because $\hat{\zeta}_2 = CE_p(\hat{\zeta})$. So, $st = st_e$.

b.2) $\hat{\zeta}$ is not an entry

By *IH* for $\hat{\zeta}_2 \rightsquigarrow^* \hat{\zeta}$ we get $st \equiv tf :: st_e$ and $tf(k_2) = \hat{c}$. This is the desired result for $\hat{\zeta}_e \rightsquigarrow^* \hat{\zeta}$. \square

Lemma 5 (Local simulation).

If $\hat{\zeta} \rightsquigarrow \hat{\zeta}'$ and $\text{succ}(|\hat{\zeta}|_{al}) \neq \emptyset$, then $|\hat{\zeta}'|_{al} \in \text{succ}(|\hat{\zeta}|_{al})$.

Proof By cases on the abstract transition.

We only show the lemma for the \widehat{UEval} to \widehat{UApply} transition, the remaining cases are similar.

$$(\llbracket (f e q)^l \rrbracket, st, h) \rightsquigarrow (f', \hat{d}, \hat{c}, st', h)$$

$$f' \in \mathcal{A}_u(f, st, h)$$

$$\hat{d} = \mathcal{A}_u(e, st, h)$$

$$\hat{c} = \mathcal{A}_k(q, st)$$

$$st' = \begin{cases} \text{pop}(st) & \text{Var}_?(q) \\ st & \text{Lam}_?(q) \wedge (H_?(f) \vee \text{Lam}_?(f)) \\ \text{setTop}([f \mapsto \{f'\}], st) & \text{Lam}_?(q) \wedge S_?(f) \end{cases}$$

A \widehat{UEval} state has a successor only when its stack is not empty, so $st \equiv tf :: st''$.

Thus, $|st|_{al} = \{(v, tf(v)) : v \in \text{dom}(tf) \wedge U\text{Var}_?(v)\}$.

Then, $|\hat{\zeta}|_{al} = (\llbracket (f e q)^l \rrbracket, |st|_{al}, h)$. Also, $|\hat{\zeta}'|_{al} = (f', \hat{d}, h)$.

It suffices to show that

$$f' \in \hat{\mathcal{A}}_u(f, |st|_{al}, h) \text{ and } \hat{d} = \hat{\mathcal{A}}_u(e, |st|_{al}, h);$$

but these hold because $\hat{\mathcal{A}}_u(v, |st|_{al}, h) = \mathcal{A}_u(v, st, h)$ is true for any v (*uvar* or *ulam*). \square

Lemma 6 (Converse of Local Simulation).

If $\tilde{\zeta} \rightsquigarrow \tilde{\zeta}'$ then, for any $\hat{\zeta}$ such that $\tilde{\zeta} = |\hat{\zeta}|_{al}$, there exists a state $\hat{\zeta}'$ such that $\hat{\zeta} \rightsquigarrow \hat{\zeta}'$ and $\tilde{\zeta}' = |\hat{\zeta}'|_{al}$ \square

Lemma 7 (Path decomposition). Let $p \equiv \hat{\zeta}_e \rightsquigarrow^* \hat{\zeta}$ be push monotonic and $\hat{\zeta}_e = (\llbracket (\lambda_l (u k) \text{ call}) \rrbracket, \hat{d}, \hat{c}, st_e, h_e)$.

- if $\hat{\zeta}$ is a \widehat{CAppl} y of the form $(\hat{c}, \dots, st_e, \dots)$ then $CE_p(\hat{\zeta})$ is not defined.
- Otherwise,

1. $CE_p(\hat{\zeta})$ is defined, i.e., $p \equiv \hat{\zeta}_e \rightsquigarrow^* \hat{\zeta}_1 \rightsquigarrow^* \hat{\zeta}$, where $\hat{\zeta}_1 = CE_p(\hat{\zeta})$.
2. Regarding the set $CE_p^*(\hat{\zeta})$, p can be in one of four forms
 - (a) $p \equiv \hat{\zeta}_e \rightsquigarrow^* \hat{\zeta}$ where $\hat{\zeta}_e = CE_p(\hat{\zeta})$ and $CE_p^*(\hat{\zeta}) = \{\hat{\zeta}_e\}$
 - (b) $p \equiv e_1 \rightsquigarrow^+ c_1 \rightsquigarrow \dots \rightsquigarrow e_k \rightsquigarrow^+ c_k \rightsquigarrow \hat{\zeta}_1 \rightsquigarrow^* \hat{\zeta}$, $k > 0$, where e_i s are entries, c_i s are tail calls, $e_1 = \hat{\zeta}_e$, $e_i = CE_p(c_i)$, $\hat{\zeta}_1 = CE_p(\hat{\zeta})$ and $CE_p^*(\hat{\zeta}) = \{e_1, \dots, e_k, \hat{\zeta}_1\}$
 - (c) $p \equiv \hat{\zeta}_e \rightsquigarrow^+ c \rightsquigarrow \hat{\zeta}_1 \rightsquigarrow^* \hat{\zeta}$ where c is a call, $\hat{\zeta}_1 = CE_p(\hat{\zeta})$ and $CE_p^*(\hat{\zeta}) = \{\hat{\zeta}_1\}$
 - (d) $p \equiv \hat{\zeta}_e \rightsquigarrow^+ c \rightsquigarrow e_1 \rightsquigarrow^+ c_1 \rightsquigarrow \dots \rightsquigarrow e_k \rightsquigarrow^+ c_k \rightsquigarrow \hat{\zeta}_1 \rightsquigarrow^* \hat{\zeta}$, $k > 0$, where c is a call, e_i s are entries, c_i s are tail calls, $e_i = CE_p(c_i)$, $\hat{\zeta}_1 = CE_p(\hat{\zeta})$ and $CE_p^*(\hat{\zeta}) = \{e_1, \dots, e_k, \hat{\zeta}_1\}$

Proof By induction on the length of p .

Basecase: $\hat{\zeta}_e \rightsquigarrow^0 \hat{\zeta}_e$

Then, $\hat{\zeta} = \hat{\zeta}_e \Rightarrow \hat{\zeta}_e = CE_p(\hat{\zeta}) \Rightarrow CE_p^*(\hat{\zeta}) = \{\hat{\zeta}_e\} \Rightarrow (2a)$ holds

Inductive step: $\hat{\zeta}_e \rightsquigarrow^* \zeta' \rightsquigarrow \hat{\zeta}$

Cases on ζ' :

a) ζ' is a Call

Then, $\hat{\zeta}$ is an entry so $CE_p(\hat{\zeta}) = \hat{\zeta}$.

Also, $CE_p^*(\hat{\zeta}) = \{\hat{\zeta}\}$ so (2c) holds.

b) ζ' is a Tail Call

Then, $\hat{\zeta}$ is an entry so $CE_p(\hat{\zeta}) = \hat{\zeta}$.

To show (2), we take cases on whether (2a), (2b), (2c) or (2d) holds for ζ' .

b.1) (2a) holds for ζ' , i.e.

$p \equiv \hat{\zeta}_e \rightsquigarrow^* \zeta' \rightsquigarrow \hat{\zeta}$ where $\hat{\zeta}_e = CE_p(\zeta')$ and $CE_p^*(\zeta') = \{\hat{\zeta}_e\}$. By the second branch of the definition of CE_p^* , $CE_p^*(\zeta') \subseteq CE_p^*(\hat{\zeta})$. Hence, $CE_p^*(\hat{\zeta}) = \{\hat{\zeta}_e, \hat{\zeta}\}$, which implies that (2b) holds for $\hat{\zeta}$.

b.2) (2b) holds for ζ'

By a similar argument, we find that (2b) holds for $\hat{\zeta}$.

b.3) (2c) holds for ζ'

By a similar argument, we find that (2d) holds for $\hat{\zeta}$.

b.4) (2d) holds for ζ'

By a similar argument, we find that (2d) holds for $\hat{\zeta}$.

c) ζ' is a $\widehat{CApply} \equiv (\hat{c}, \dots, st_e, \dots)$

Then, in the transition $\zeta' \rightsquigarrow \hat{\zeta}$ we do *setTop* on st_e , which means that p isn't push monotonic. Thus, this case can't arise.

d) ζ' is an inner \widehat{CEval} or a $\widehat{CApply} \neq (\hat{c}, \dots, st_e, \dots)$

By IH, $p \equiv \hat{I}(pr) \rightsquigarrow^* \hat{\zeta}_1 \rightsquigarrow^+ \zeta' \rightsquigarrow \hat{\zeta}$, where $\hat{\zeta}_1 = CE_p(\zeta')$.

By the second branch of the definition of CE_p , $\hat{\zeta}_1 = CE_p(\hat{\zeta})$.

To show (2), we take cases on whether (2a), (2b), (2c) or (2d) holds for ζ' .

The reasoning is the same as in case (b).

e) ζ' is a \widehat{CEval} exit

By IH, $p \equiv \hat{\zeta}_e \rightsquigarrow^* \hat{\zeta}_1 \rightsquigarrow^+ \zeta' \rightsquigarrow \hat{\zeta}$, where $\hat{\zeta}_1 = CE_p(\zeta')$.

Cases on (2a), (2b), (2c) or (2d) for ζ' .

e.1) (2a) holds for ζ' , i.e.

$p \equiv \hat{\zeta}_e \rightsquigarrow^+ \zeta' \rightsquigarrow \hat{\zeta}$ where $\hat{\zeta}_e = CE_p(\zeta')$.

By lemma 4, the stack of ζ' is of the form $tf :: st_e$ and $tf(k) = \hat{c}$. Thus, $\hat{\zeta}' \equiv (\hat{c}, \dots, st_e, \dots)$. The only way for $CE_p(\hat{\zeta})$ to exist is by the third

branch of the definition of CE_p , since ζ' is a \widehat{CEval} exit. But there is no call leading to $\hat{\zeta}_e$, thus $CE_p(\hat{\zeta})$ can't exist.

Similarly when (2b) holds for ζ' .

e.2) (2c) holds for ζ' , i.e.

$p \equiv \hat{\zeta}_e \rightsquigarrow^+ c \rightsquigarrow \hat{\zeta}_1 \rightsquigarrow^+ \zeta' \rightsquigarrow \hat{\zeta}$ where c is a call and $\hat{\zeta}_1 = CE_p(\zeta')$.

By IH, $CE_p(c)$ exists so p can be written $p \equiv \hat{\zeta}_e \rightsquigarrow^* \hat{\zeta}_2 \rightsquigarrow^+ c \rightsquigarrow \hat{\zeta}_1 \rightsquigarrow^+ \zeta' \rightsquigarrow \hat{\zeta}$ where $\hat{\zeta}_2 = CE_p(c)$. Then, by the third branch of the definition of CE_p , $CE_p(\hat{\zeta}) = CE_p(c) = \hat{\zeta}_2$.

To show (2) for $\hat{\zeta}$ we work as in the previous cases.

f) $\hat{\zeta}'$ is an Entry

This case is simple. \square

Lemma 8 (Stack irrelevance).

Let $p \equiv \hat{\zeta}_1 \rightsquigarrow \hat{\zeta}_2 \rightsquigarrow \dots \rightsquigarrow \hat{\zeta}_n$ be push monotonic, where $\hat{\zeta}_1 = (\text{ulam}, \hat{d}, \hat{c}, st_e, h_e)$. Also, $\hat{\zeta}_n$ is not a \widehat{CAppl} y of the form $(\hat{c}, \dots, st_e, \dots)$. By property 1, the stack of each $\hat{\zeta}_i$ is of the form $\text{append}(st_i, st_e)$.

For an arbitrary stack st' and continuation \hat{c}' , consider the sequence p' of states $\hat{\zeta}'_1 \hat{\zeta}'_2 \dots \hat{\zeta}'_n$ where each $\hat{\zeta}'_i$ is produced by $\hat{\zeta}_i$ as follows:

- if $\hat{\zeta}_i$ is an entry with stack st_e then replace the continuation argument with \hat{c}' and the stack with st' .
- if st_e is a proper suffix of the stack of $\hat{\zeta}_i$ then the latter is of the form $\text{append}(st'_i, \langle fr_i \rangle, st_e)$ for some stack st'_i . Replace st_e with st' and bind the continuation variable in fr_i to \hat{c}' .

(Note: the map isn't total, but it should be defined for all states in p .)

Then,

- for any two states $\hat{\zeta}'_i$ and $\hat{\zeta}'_{i+1}$ in p' , $(\hat{\zeta}'_i, \hat{\zeta}'_{i+1}) \in \rightsquigarrow$
- the path p' is push monotonic

Proof By induction on the length of p .

The basecase is simple.

Inductive step: $p = \hat{\zeta}_1 \rightsquigarrow^* \hat{\zeta}_{n-1} \rightsquigarrow \hat{\zeta}_n$

By IH, the transitions in the path $\hat{\zeta}'_1 \rightsquigarrow^* \hat{\zeta}'_{n-1}$ are valid with respect to the abstract semantics and the path is push monotonic. We must show that $(\hat{\zeta}'_{n-1}, \hat{\zeta}'_n) \in \rightsquigarrow$ and that $\hat{\zeta}'_1 \rightsquigarrow^* \hat{\zeta}'_n$ is push monotonic.

Cases on $\hat{\zeta}_{n-1}$:

1. $\hat{\zeta}_{n-1}$ is a \widehat{UEval} , of the form $(\llbracket (f e q)^l \rrbracket, st, h)$

By lemma 7, $CE_p(\hat{\zeta}_{n-1})$ is defined and p can be in one of four forms. We consider only the first case, the rest are similar.

Let $p \equiv \hat{\zeta}_1 \rightsquigarrow^+ \hat{\zeta}_{n-1} \rightsquigarrow \hat{\zeta}_n$ where $\hat{\zeta}_1 = CE_p(\hat{\zeta}_{n-1})$.

By lemma 4, st is of the form $tf :: st_e$ and the continuation variable in tf (call it k) is bound to \hat{c} .

- (a) q is a variable

By the abstract semantics, $\hat{\zeta}_n$ is $(\text{ulam}_n, \hat{d}_n, \hat{c}, st_e, h)$. Also, the state $\hat{\zeta}'_{n-1}$ is $(\llbracket (f e q)^l \rrbracket, tf[k \mapsto \hat{c}'] :: st', h)$, and it transitions to $(\text{ulam}_n, \hat{d}_n, \hat{c}', st', h)$ which is $\hat{\zeta}'_n$.

- (b) q is a lambda and f is a stack reference

Then, $\hat{\zeta}_n$ is $(\text{ulam}_n, \hat{d}_n, q, tf[f \mapsto \{\text{ulam}_n\}] :: st_e, h)$. Also, the state $\hat{\zeta}'_{n-1}$ is $(\llbracket (f e q)^l \rrbracket, tf[k \mapsto \hat{c}'] :: st', h)$, and it transitions to $(\text{ulam}_n, \hat{d}_n, q, tf[k \mapsto \hat{c}'] [f \mapsto \{\text{ulam}_n\}] :: st', h)$ which is $\hat{\zeta}'_n$.

(c) q is a lambda and f is a heap reference

Similarly.

2. $\hat{\zeta}_{n-1}$ is a \widehat{CEval} exit

By lemma 7, $CE_p(\hat{\zeta}_{n-1})$ is defined and p can be in one of four forms.

(a) $p \equiv \hat{\zeta}_1 \rightsquigarrow^+ \hat{\zeta}_{n-1} \rightsquigarrow \hat{\zeta}_n$ where $\hat{\zeta}_1 = CE_p(\hat{\zeta}_{n-1})$

Then, by lemma 4 and the abstract semantics, it is easy to see that $\hat{\zeta}_n$ is of the form $(\hat{c}, \dots, st_e, \dots)$. Thus, this case isn't possible.

Similarly when $\hat{\zeta}_1 \neq CE_p(\hat{\zeta}_{n-1})$ but is in $CE_p^*(\hat{\zeta}_{n-1})$.

(b) $p \equiv \hat{\zeta}_1 \rightsquigarrow^+ c \rightsquigarrow \hat{\zeta}'_e \rightsquigarrow^+ \hat{\zeta}_{n-1} \rightsquigarrow \hat{\zeta}$ where $\hat{\zeta}'_e = CE_p(\hat{\zeta}_{n-1})$ and c is a call:

Then, $CE_p(c)$ is defined and its stack has st_e as a suffix. Hence, by lemma 4, the stack of c is bigger than st_e by at least a frame. Since the stack of $\hat{\zeta}'_e$ has the same size as the stack of c , the stack of $\hat{\zeta}_{n-1}$ is bigger than st_e by at least two frames. By lemma 2 we get the desired.

Similarly when $\hat{\zeta}'_e \neq CE_p(\hat{\zeta}_{n-1})$ but is in $CE_p^*(\hat{\zeta}_{n-1})$.

3. $\hat{\zeta}_{n-1}$ is an inner \widehat{CEval}

Similarly to the previous cases.

4. $\hat{\zeta}_{n-1}$ is a \widehat{UApply}

Lemma 7 gives the same four cases. We only consider one, the rest are similar.

Let $p \equiv \hat{\zeta}_1 \rightsquigarrow^+ c \rightsquigarrow \hat{\zeta}_{n-1} \rightsquigarrow \hat{\zeta}_n$ where c is a call.

Then, $CE_p(c)$ is defined and its stack has st_e as a suffix. Hence, by lemma 4, the stack of c is bigger than st_e by at least a frame. Since the stack of $\hat{\zeta}_{n-1}$ has the same size as the stack of c , we don't change the continuation argument in $\hat{\zeta}'_{n-1}$. By lemma 3 we get the desired.

5. $\hat{\zeta}_{n-1}$ is a \widehat{CAppl}

Similarly to the previous cases. □

Theorem 5 (Soundness).

If $p \equiv \hat{\mathcal{I}}(pr) \rightsquigarrow^* \hat{\zeta}$ then, after summarization:

- if $\hat{\zeta}$ is not a final state then $(|CE_p(\hat{\zeta})|_{al}, |\hat{\zeta}|_{al}) \in Seen$
- if $\hat{\zeta}$ is a final state then $|\hat{\zeta}|_{al} \in Final$
- if $\hat{\zeta}$ is a \widehat{CEval} exit and $\hat{\zeta}' \in CE_p^*(\hat{\zeta})$ then $(|\hat{\zeta}'|_{al}, |\hat{\zeta}|_{al}) \in Seen$

Proof By induction on the length of p .

Basecase: $\hat{\mathcal{I}}(pr) \rightsquigarrow^0 \hat{\mathcal{I}}(pr)$

Then, $(\hat{\mathcal{I}}(pr), \hat{\mathcal{I}}(pr)) \in Seen$.

Inductive step: $\hat{\mathcal{I}}(pr) \rightsquigarrow^* \hat{\zeta}' \rightsquigarrow \hat{\zeta}$

Cases on $\hat{\zeta}$:

a) $\hat{\zeta}$ is an Entry

Then, $CE_p(\hat{\zeta}) = \hat{\zeta}$. Also, $\hat{\zeta}'$ is a call or a tail call.

By lemma 7, $p \equiv \hat{\mathcal{I}}(pr) \rightsquigarrow^* \hat{\zeta}_1 \rightsquigarrow^+ \hat{\zeta}' \rightsquigarrow \hat{\zeta}$, where $\hat{\zeta}_1 = CE_p(\hat{\zeta}')$.

By IH, $(|\hat{\zeta}_1|_{al}, |\hat{\zeta}'|_{al}) \in Seen$ which means that it has been entered in W and examined. By lemma 5, $|\hat{\zeta}|_{al} \in succ(|\hat{\zeta}'|_{al})$ so in line 10 or 22 $(|\hat{\zeta}|_{al}, |\hat{\zeta}|_{al})$ will be propagated.

b) $\hat{\zeta}$ is a \widehat{CAppl} y but not a final state

Then, $\hat{\zeta} = (\llbracket (\lambda_\gamma(u) \text{ call}) \rrbracket, \hat{d}, st, h)$

and $\hat{\zeta}' = (\llbracket (q e)^{\gamma'} \rrbracket, st', h)$.

b.1) $Lam_\gamma(q)$, i.e. $\hat{\zeta}'$ is an inner \widehat{CEval}

This case is simple.

b.2) $Var_\gamma(q)$, i.e. $\hat{\zeta}'$ is a \widehat{CEval} exit

The path $\hat{I}(pr) \rightsquigarrow^* \hat{\zeta}'$ satisfies part 2 of lemma 7. It can't satisfy cases 2a or 2b because $\hat{\zeta}$ would be a final state by lemma 4. Thus, it satisfies 2c or 2d. Then, the path is of the form

$$p \equiv \hat{I}(pr) \rightsquigarrow^* \hat{\zeta}_1 \rightsquigarrow^+ \hat{\zeta}_2 \rightsquigarrow \hat{\zeta}_3 \rightsquigarrow^+ \hat{\zeta}' \rightsquigarrow \hat{\zeta}$$

where $\hat{\zeta}_2$ is a call, $\hat{\zeta}_1 = CE_p(\hat{\zeta}_2)$ and $\hat{\zeta}_3 \in CE_p^*(\hat{\zeta}')$.

Note that by the third branch of the definition of CE_p , $\hat{\zeta}_1 = CE_p(\hat{\zeta})$. We must show that $(|\hat{\zeta}_1|_{al}, |\hat{\zeta}|_{al}) \in Seen$.

The state $\hat{\zeta}_1$ is an entry of the form

$$\hat{\zeta}_1 = (\llbracket (\lambda_{l_1}(u_1 k_1) \text{ call}_1) \rrbracket, \hat{d}_1, \hat{c}_1, st_1, h_1)$$

The state $\hat{\zeta}_2$ is a call of the form

$$\hat{\zeta}_2 = (\llbracket (f_2 e_2 q_2)^{l_2} \rrbracket, st_2, h_2), \text{ where } q_2 \text{ is a } clam.$$

Lemma 4 for $\hat{\zeta}_1 \rightsquigarrow^+ \hat{\zeta}_2$ gives $st_2 \equiv tf_2 :: st_1$.

By the abstract semantics for $\hat{\zeta}_2 \rightsquigarrow \hat{\zeta}_3$, we get:

$$\hat{\zeta}_3 = (ulam, \hat{d}_3, q_2, st_3, h_2), \text{ where}$$

either $st_3 = st_2$, if $(Lam_\gamma(f_2) \vee H_\gamma(f_2))$ holds,

or $st_3 = setTop(\llbracket f_2 \mapsto \{ulam\} \rrbracket, st_2)$, if $S_\gamma(f_2)$ holds.

i.e. $st_3 = tf_3 :: st_1$, and

$$tf_3 = \begin{cases} tf_2 & Lam_\gamma(f_2) \vee H_\gamma(f_2) \\ tf_2 \llbracket f_2 \mapsto \{ulam\} \rrbracket & S_\gamma(f_2) \end{cases}$$

By lemma 4 for $\hat{\zeta}_3 \rightsquigarrow^+ \hat{\zeta}'$, we get $st' = tf' :: st_3$ and $tf'(q) = q_2$.

Then, by the abstract semantics for $\hat{\zeta}' \rightsquigarrow \hat{\zeta}$,

$$q_2 = \llbracket (\lambda_\gamma(u) \text{ call}) \rrbracket, st = st_3, \text{ and } \hat{d} = \mathcal{A}_u(e, st', h).$$

The above information will become useful when dealing with the local counterparts of the aforementioned states.

By *IH*, $(|\hat{\zeta}_3|_{al}, |\hat{\zeta}'|_{al})$ was entered in W (at line 25) and later examined at line 13. Note that $\hat{\zeta}_3 \neq \hat{I}(pr)$ because $\hat{\zeta}_2$ is between them, therefore **Final** will not be called at line 15.

Also by *IH*, $(|\hat{\zeta}_1|_{al}, |\hat{\zeta}_2|_{al})$ was entered in W and later examined. Lemma 5 implies that $|\hat{\zeta}_3|_{al} \in succ(|\hat{\zeta}_2|_{al})$ so $(|\hat{\zeta}_1|_{al}, |\hat{\zeta}_2|_{al}, |\hat{\zeta}_3|_{al})$ will go in *Callers*.

We take cases on whether $(|\hat{\zeta}_3|_{al}, |\hat{\zeta}'|_{al})$ or $(|\hat{\zeta}_1|_{al}, |\hat{\zeta}_2|_{al})$ was examined first by the algorithm.

b.2.1) $(|\hat{\zeta}_1|_{al}, |\hat{\zeta}_2|_{al})$ was examined first

Then, when $(|\hat{\zeta}_3|_{al}, |\hat{\zeta}'|_{al})$ is examined,

$(|\hat{\zeta}_1|_{al}, |\hat{\zeta}_2|_{al}, |\hat{\zeta}_3|_{al})$ is in *Callers*.

Therefore, at line 18 we call

Update $(|\hat{\zeta}_1|_{al}, |\hat{\zeta}_2|_{al}, |\hat{\zeta}_3|_{al}, |\hat{\zeta}'|_{al})$.

By applying $|\cdot|_{al}$ to the abstract states we get

$$|\hat{\zeta}_1|_{al} = (\llbracket (\lambda_{l_1} (u_1 k_1) call_1) \rrbracket, \hat{d}_1, h_1)$$

$$|\hat{\zeta}_2|_{al} = (\llbracket (f_2 e_2 q_2)^{l_2} \rrbracket, tf_2, h_2),$$

where $q_2 = \llbracket (\lambda_\gamma (u) call) \rrbracket$.

$$|\hat{\zeta}_3|_{al} = (ulam, \hat{d}_3, h_2)$$

$$|\zeta'|_{al} = (\llbracket (qe)^\gamma \rrbracket, tf', h),$$

where $tf'(q) = \llbracket (\lambda_\gamma (u) call) \rrbracket$.

By looking at the code for **Update**, we see that the return value is $\hat{\mathcal{A}}_u(e, tf', h) = \mathcal{A}_u(e, st', h) = \hat{d}$. The frame of the return state is

$$\begin{cases} tf_2 & Lam_\gamma(f_2) \vee H_\gamma(f_2) \\ tf_2[f_2 \mapsto \{ulam\}] & S_\gamma(f_2) \end{cases}$$

which is equal to tf_3 . The heap at the return state is h . Last, the continuation we are returning to is $\llbracket (\lambda_\gamma (u) call) \rrbracket$. Thus, the return state $\tilde{\zeta}$ is equal to $|\hat{\zeta}|_{al}$, and we call **Propagate** $(|\hat{\zeta}_1|_{al}, |\hat{\zeta}|_{al})$, so $(|\hat{\zeta}_1|_{al}, |\hat{\zeta}|_{al})$ will go in *Seen*.

b.2.2) $(|\hat{\zeta}_3|_{al}, |\zeta'|_{al})$ was examined first

Then, when $(|\hat{\zeta}_1|_{al}, |\hat{\zeta}_2|_{al})$ is examined, $(|\hat{\zeta}_3|_{al}, |\zeta'|_{al})$ is in *Summary*, and at line 12 we call

Update $(|\hat{\zeta}_1|_{al}, |\hat{\zeta}_2|_{al}, |\hat{\zeta}_3|_{al}, |\zeta'|_{al})$.

Proceed as above.

c) $\hat{\zeta}$ is a final state

Then, $\hat{\zeta} = (halt, \hat{d}, \langle \rangle, h)$.

We must show that $|\hat{\zeta}|_{al}$ will be in *Final* after the execution of the summarization algorithm.

By the abstract semantics for $\zeta' \rightsquigarrow \hat{\zeta}$,

$$\zeta' = (\llbracket (ke)^\gamma \rrbracket, st', h),$$

where $st' = tf' :: \langle \rangle$, $tf'(k) = halt$, and $\hat{d} = \mathcal{A}_u(e, st', h)$.

By *IH* for $\hat{\mathcal{I}}(pr) \rightsquigarrow^* \zeta'$, we know that $(|\hat{\mathcal{I}}(pr)|_{al}, |\zeta'|_{al})$ was entered in W and *Summary* sometime during the algorithm. When it was examined, the test at line 14 was true so we called **Final** $(|\zeta'|_{al})$. Hence, we insert $\tilde{\zeta} = (halt, \hat{\mathcal{A}}_u(e, tf', h), \langle \rangle, h)$ in *Final*. But, $\hat{\mathcal{A}}_u(e, tf', h) = \mathcal{A}_u(e, st', h) = \hat{d}$, hence $\tilde{\zeta} = |\hat{\zeta}|_{al}$.

d) $\hat{\zeta}$ is a \widehat{CEval} exit

By lemma 7 for $\hat{\mathcal{I}}(pr) \rightsquigarrow^* \zeta'$, $p \equiv \hat{\mathcal{I}}(pr) \rightsquigarrow^* \hat{\zeta}_1 \rightsquigarrow^* \zeta' \rightsquigarrow \hat{\zeta}$, where $\hat{\zeta}_1 = CE_p(\zeta')$. But ζ' is not a \widehat{CEval} exit (it is an *Apply* state), so by the second branch of the definition of CE_p we get $\hat{\zeta}_1 = CE_p(\hat{\zeta})$.

By *IH*, $(|\hat{\zeta}_1|_{al}, |\zeta'|_{al})$ is entered in *Seen* and W ; and examined at line 6. By lemma 5, $|\zeta'|_{al} \in succ(|\zeta'|_{al})$ so $(|\hat{\zeta}_1|_{al}, |\zeta'|_{al})$ will be propagated (line 7) and entered in *Seen* (line 25).

We need to show that for every $\zeta'' \in CE_p^*(\hat{\zeta})$, $(|\zeta''|_{al}, |\hat{\zeta}|_{al})$ will be inserted in *Seen*. The path $\hat{\mathcal{I}}(pr) \rightsquigarrow^* \zeta'$ satisfies part 2 of lemma 7; proceed by cases:

d.1) $\hat{\mathcal{I}}(pr) \rightsquigarrow^* \zeta'$ satisfies 2a

Then, $\hat{\zeta}_1 = \hat{\mathcal{I}}(pr)$ and $p \equiv \hat{\zeta}_1 \rightsquigarrow^* \zeta' \rightsquigarrow \hat{\zeta}$ and $CE_p^*(\hat{\zeta}) = \{\hat{\zeta}_1\}$. But we've shown that $(|\hat{\zeta}_1|_{al}, |\hat{\zeta}|_{al})$ is entered in *Seen*.

- d.2) $\hat{I}(pr) \rightsquigarrow^* \zeta'$ satisfies 2b
Then, $p \equiv e_1 \rightsquigarrow^+ c_1 \rightsquigarrow \dots \rightsquigarrow e_k \rightsquigarrow^+ c_k \rightsquigarrow \hat{\zeta}_1 \rightsquigarrow^* \zeta' \rightsquigarrow \hat{\zeta}$, where $e_1 = \hat{I}(pr)$, e_i s are entries, c_i s are tail calls, $e_i = CE_p(c_i)$, $CE_p^*(\zeta') = \{e_1, \dots, e_k, \hat{\zeta}_1\}$.
Hence, $CE_p^*(\hat{\zeta}) = \{e_1, \dots, e_k, \hat{\zeta}_1\}$. To show that $(|e_k|_{al}, |\hat{\zeta}|_{al})$ is entered in *Seen*, we proceed by cases on whether $(|e_k|_{al}, |c_k|_{al})$ or $(|\hat{\zeta}_1|_{al}, |\hat{\zeta}|_{al})$ was examined first by the algorithm.
- d.2.1) $(|e_k|_{al}, |c_k|_{al})$ was examined first
By lemma 5, $|\hat{\zeta}_1|_{al}$ is in $succ(|c_k|_{al})$, therefore $(|e_k|_{al}, |c_k|_{al}, |\hat{\zeta}_1|_{al})$ will go in *TCallers*. Then, when $(|\hat{\zeta}_1|_{al}, |\hat{\zeta}|_{al})$ is examined, in line 19 we will call **Propagate** $(|e_k|_{al}, |\hat{\zeta}|_{al})$, so $(|e_k|_{al}, |\hat{\zeta}|_{al})$ will go in *Seen*.
- d.2.2) $(|\hat{\zeta}_1|_{al}, |\hat{\zeta}|_{al})$ was examined first
When $(|e_k|_{al}, |c_k|_{al})$ is examined, $(|\hat{\zeta}_1|_{al}, |\hat{\zeta}|_{al})$ will be in *Summary*, and by lemma 5 we know $|\hat{\zeta}_1|_{al} \in succ(|c_k|_{al})$. Thus, in line 24 we will call **Propagate** which will insert $(|e_k|_{al}, |\hat{\zeta}|_{al})$ in *Seen*.
By repeating this process $k-1$ times, we can show that all edges $(|e_i|_{al}, |\hat{\zeta}|_{al})$ go in *Seen*.
- d.3) $\hat{I}(pr) \rightsquigarrow^* \zeta'$ satisfies 2c or 2d
These cases are similar to the previous cases. The only difference is that now $\hat{I}(pr)$ is not in $CE_p^*(\zeta')$ (which doesn't change the proof).
- e) $\hat{\zeta}$ is a Tail Call (thus an exit)
By lemma 7 for $\hat{I}(pr) \rightsquigarrow^* \zeta'$, $p \equiv \hat{I}(pr) \rightsquigarrow^* \hat{\zeta}_1 \rightsquigarrow^* \zeta' \rightsquigarrow \hat{\zeta}$, where $\hat{\zeta}_1 = CE_p(\zeta')$. But ζ' is not a \widehat{CEval} exit (it is an \widehat{Apply} state), so by the second branch of the definition of CE_p we get $\hat{\zeta}_1 = CE_p(\hat{\zeta})$.
By IH, $(|\hat{\zeta}_1|_{al}, |\zeta'|_{al})$ is entered in *Seen* and *W*; and examined at line 6. By lemma 5, $|\hat{\zeta}|_{al} \in succ(|\zeta'|_{al})$ so $(|\hat{\zeta}_1|_{al}, |\hat{\zeta}|_{al})$ will be propagated (line 7) and entered in *Seen* (line 25).
- f) $\hat{\zeta}$ is an inner \widehat{CEval}
This case is simple.
- g) $\hat{\zeta}$ is a Call
This case is simple. □

Theorem 6 (Completeness).

After summarization:

- For each $(\tilde{\zeta}_1, \tilde{\zeta}_2)$ in *Seen*, there exist $\hat{\zeta}_1, \hat{\zeta}_2$ and p such that $p \equiv \hat{I}(pr) \rightsquigarrow^* \hat{\zeta}_1 \rightsquigarrow^* \hat{\zeta}_2$ and $\tilde{\zeta}_1 = |\hat{\zeta}_1|_{al}$ and $\tilde{\zeta}_2 = |\hat{\zeta}_2|_{al}$ and $\hat{\zeta}_1 \in CE_p^*(\hat{\zeta}_2)$
- For each $\tilde{\zeta}$ in *Final*, there exist $\hat{\zeta}$ and p such that $p \equiv \hat{I}(pr) \rightsquigarrow^+ \hat{\zeta}$ and $\tilde{\zeta} = |\hat{\zeta}|_{al}$ and $\hat{\zeta}$ is a final state.

Proof By induction on the number of iterations. We prove that the algorithm maintains the following properties for *Seen* and *Final*.

1. For each $(\tilde{\zeta}_1, \tilde{\zeta}_2)$ in *Seen*, there exist $\hat{\zeta}_1, \hat{\zeta}_2$ and p such that $p \equiv \hat{I}(pr) \rightsquigarrow^* \hat{\zeta}_1 \rightsquigarrow^* \hat{\zeta}_2$ and $\tilde{\zeta}_1 = |\hat{\zeta}_1|_{al}$ and $\tilde{\zeta}_2 = |\hat{\zeta}_2|_{al}$ and

- if $\tilde{\zeta}_2$ is a \widetilde{CEval} exit then $\hat{\zeta}_1 \in CE_p^*(\hat{\zeta}_2)$,
otherwise $\hat{\zeta}_1 = CE_p(\hat{\zeta}_2)$
- 2. For each $\tilde{\zeta}$ in *Final*, there exist $\hat{\zeta}$ and p such that
 $p \equiv \hat{\mathcal{I}}(pr) \rightsquigarrow^+ \hat{\zeta}$ and
 $\tilde{\zeta} = |\hat{\zeta}|_{al}$ and
 $\hat{\zeta}$ is a final state.

Initially, we must show that the properties hold before the first iteration (at the beginning of the algorithm): *Final* is empty and W contains just $(\tilde{\mathcal{I}}(pr), \tilde{\mathcal{I}}(pr))$, for which property 1 holds.

Now the inductive step: at the beginning of each iteration, we remove an edge $(\tilde{\zeta}_1, \tilde{\zeta}_2)$ from W . We assume that the properties hold at that point. We must show that, after we process the edge, the new elements of *Seen* and *Final* satisfy the properties.

- $\tilde{\zeta}_2$ is an entry, a \widetilde{CAppl} y or an inner \widetilde{CEval}
 $(\tilde{\zeta}_1, \tilde{\zeta}_2)$ is in *Seen*, so by *IH*
 $\exists \hat{\zeta}_1, \hat{\zeta}_2, p. p \equiv \hat{\mathcal{I}}(pr) \rightsquigarrow^* \hat{\zeta}_1 \rightsquigarrow^* \hat{\zeta}_2 \wedge \tilde{\zeta}_1 = |\hat{\zeta}_1|_{al} \wedge \tilde{\zeta}_2 = |\hat{\zeta}_2|_{al} \wedge \hat{\zeta}_1 = CE_p(\hat{\zeta}_2)$
For each $\tilde{\zeta}_3$ in $\text{succ}(\tilde{\zeta}_2)$, $(\tilde{\zeta}_1, \tilde{\zeta}_3)$ will be propagated.
If $(\tilde{\zeta}_1, \tilde{\zeta}_3)$ is already in *Seen* then property 1 holds by *IH* (in the following cases, we won't repeat this argument and will assume that the insertion in *Seen* happens now).
Otherwise, we insert the edge at this iteration, at line 25. By lemma 6,
 $\exists \hat{\zeta}_3. \tilde{\zeta}_3 = |\hat{\zeta}_3|_{al} \wedge \hat{\zeta}_2 \rightsquigarrow \hat{\zeta}_3$
By the second branch of the definition of CE_p , $\hat{\zeta}_1 = CE_p(\hat{\zeta}_3)$
- $\tilde{\zeta}_2$ is a call
Let $\tilde{\zeta}_1 = (\llbracket (\lambda_1(u_1 k_1) call_1) \rrbracket, \hat{d}_1, h_1)$
and $\tilde{\zeta}_2 = (\llbracket (f_2 e_2 (\lambda_2(u_2) call_2)) \rrbracket, tf_2, h_2)$
Also, assume $S_?(f_2)$ (the other cases are simpler).
 $(\tilde{\zeta}_1, \tilde{\zeta}_2)$ is in *Seen*, so by *IH*
 $\exists \hat{\zeta}_1, \hat{\zeta}_2, p. p \equiv \hat{\mathcal{I}}(pr) \rightsquigarrow^* \hat{\zeta}_1 \rightsquigarrow^+ \hat{\zeta}_2 \wedge \tilde{\zeta}_1 = |\hat{\zeta}_1|_{al} \wedge \tilde{\zeta}_2 = |\hat{\zeta}_2|_{al} \wedge \hat{\zeta}_1 = CE_p(\hat{\zeta}_2)$
Each entry $\tilde{\zeta}_3$ in $\text{succ}(\tilde{\zeta}_2)$ will be propagated. By lemma 6,
 $\exists \hat{\zeta}_3. \tilde{\zeta}_3 = |\hat{\zeta}_3|_{al} \wedge \hat{\zeta}_2 \rightsquigarrow \hat{\zeta}_3$
Since $\hat{\zeta}_3 = CE_p(\hat{\zeta}_3)$, property 1 holds for $\tilde{\zeta}_3$.
If there is no edge $(\tilde{\zeta}_3, \tilde{\zeta}_4)$ in *Summary*, we are done. Otherwise, we call **Update** $(\tilde{\zeta}_1, \tilde{\zeta}_2, \tilde{\zeta}_3, \tilde{\zeta}_4)$ and we must show that property 1 holds for the edge inserted in *Seen* by **Update**.
Let st_1 be the stack of $\hat{\zeta}_1$. By lemma 4, the stack of $\hat{\zeta}_2$ is $tf_2 :: st_1$.
Let $\tilde{\zeta}_3 = (\llbracket (\lambda_3(u_3 k_3) call_3) \rrbracket, \hat{d}_3, h_2)$
and $\tilde{\zeta}_4 = (\llbracket (k_4 e_4) \rrbracket, tf_4, h_4)$.
(Note that tf_4 contains only user bindings.)
We know $\text{Summary} \subseteq \text{Seen}$ so by *IH* for $(\tilde{\zeta}_3, \tilde{\zeta}_4)$ we get (note that $\tilde{\zeta}_4$ is a \widetilde{CEval} exit)
 $\exists \hat{\zeta}'_3, \hat{\zeta}'_4, p'. p' \equiv \hat{\mathcal{I}}(pr) \rightsquigarrow^* \hat{\zeta}'_3 \rightsquigarrow^+ \hat{\zeta}'_4 \wedge \tilde{\zeta}_3 = |\hat{\zeta}'_3|_{al} \wedge \tilde{\zeta}_4 = |\hat{\zeta}'_4|_{al} \wedge \hat{\zeta}'_3 \in CE_{p'}^*(\hat{\zeta}'_4)$
Then, $\hat{\zeta}'_3 = (\llbracket (\lambda_3(u_3 k_3) call_3) \rrbracket, \hat{d}_3, \hat{c}_3, st'_3, h_2)$
and by lemma 4, $\hat{\zeta}'_4 = (\llbracket (k_4 e_4) \rrbracket, tf_4[k_4 \mapsto \hat{c}_3] :: st'_3, h_4)$.

But the path from $\hat{\zeta}'_3$ to $\hat{\zeta}'_4$ is push monotonic, so by lemma 8 there exist states

$$\hat{\zeta}_3 = (\llbracket (\lambda_3(u_3 k_3) \text{ call}_3) \rrbracket, \hat{d}_3, \llbracket (\lambda_2(u_2) \text{ call}_2) \rrbracket, st_3, h_2)$$

where $st_3 = tf_2[f_2 \mapsto \{\llbracket (\lambda_3(u_3 k_3) \text{ call}_3) \rrbracket\}] :: st_1$,

and $\hat{\zeta}_4 = (\llbracket (k_4 e_4) \rrbracket, st_4, h_4)$

where $st_4 = tf_4[k_4 \mapsto \llbracket (\lambda_2(u_2) \text{ call}_2) \rrbracket] :: st_3$,

such that $\hat{\zeta}_3 \rightsquigarrow^+ \hat{\zeta}_4$.

Thus, the path p can be extended to

$$\hat{\mathcal{I}}(pr) \rightsquigarrow^* \hat{\zeta}_1 \rightsquigarrow^+ \hat{\zeta}_2 \rightsquigarrow \hat{\zeta}_3 \rightsquigarrow^+ \hat{\zeta}_4$$

By the abstract semantics, the successor $\hat{\zeta}$ of $\hat{\zeta}_4$ is

$$(\llbracket (\lambda_2(u_2) \text{ call}_2) \rrbracket, \mathcal{A}_u(e_4, st_4, h_4), st_3, h_4).$$

The state $\tilde{\zeta}$ produced by **Update** is

$$(\llbracket (\lambda_2(u_2) \text{ call}_2) \rrbracket, \hat{\mathcal{A}}_u(e_4, tf_4, h_4), tf, h_4)$$

where $tf = tf_2[f_2 \mapsto \{\llbracket (\lambda_3(u_3 k_3) \text{ call}_3) \rrbracket\}]$.

It is simple to see that $\tilde{\zeta} = |\hat{\zeta}|_{al}$.

– $\tilde{\zeta}_2$ is a *CEval* exit, $(\llbracket (k e) \rrbracket, tf_2, h_2)$

If $\tilde{\zeta}_1$ is $\hat{\mathcal{I}}(pr)$ then **Final**($\tilde{\zeta}_2$) is called and a local state $\tilde{\zeta}$ of the form $(halt, \hat{\mathcal{A}}_u(e, tf_2, h_2), \langle \rangle, h_2)$ goes in *Final*. We must show that property 2 holds. By *IH* for $(\tilde{\zeta}_1, \tilde{\zeta}_2)$,

$$\exists \hat{\zeta}_2, p. p \equiv \hat{\mathcal{I}}(pr) \rightsquigarrow^+ \hat{\zeta}_2 \wedge \tilde{\zeta}_2 = |\hat{\zeta}_2|_{al} \wedge \hat{\mathcal{I}}(pr) \in CE_p^*(\hat{\zeta}_2)$$

(Note that $\hat{\zeta}_1 = \hat{\mathcal{I}}(pr)$) By lemma 4, the stack st_2 of $\hat{\zeta}_2$ is $tf_2[k \mapsto halt] :: \langle \rangle$.

Hence, the successor $\hat{\zeta}$ of $\hat{\zeta}_2$ is $(halt, \mathcal{A}_u(e, st_2, h_2), \langle \rangle, h_2)$, and $\tilde{\zeta} = |\hat{\zeta}|_{al}$ holds.

If $\tilde{\zeta}_1 \neq \hat{\mathcal{I}}(pr)$, for each triple $(\tilde{\zeta}_3, \tilde{\zeta}_4, \tilde{\zeta}_1)$ in *Callers*, we call **Update**($\tilde{\zeta}_3, \tilde{\zeta}_4, \tilde{\zeta}_1, \tilde{\zeta}_2$).

Insertion in *Callers* happens only at line 11, which means that $(\tilde{\zeta}_3, \tilde{\zeta}_4)$ is in *Seen*. Thus, by *IH*

$$\exists \hat{\zeta}_3, \hat{\zeta}_4, p. p \equiv \hat{\mathcal{I}}(pr) \rightsquigarrow^* \hat{\zeta}_3 \rightsquigarrow^+ \hat{\zeta}_4 \wedge \tilde{\zeta}_3 = |\hat{\zeta}_3|_{al} \wedge \tilde{\zeta}_4 = |\hat{\zeta}_4|_{al} \wedge \hat{\zeta}_3 = CE_p(\hat{\zeta}_4)$$

Also, $\tilde{\zeta}_4 \approx \tilde{\zeta}_1$ thus by lemma 6

$$\exists \hat{\zeta}_1. \hat{\zeta}_4 \rightsquigarrow \hat{\zeta}_1 \wedge \tilde{\zeta}_1 = |\hat{\zeta}_1|_{al}$$

Using the *IH* for $(\tilde{\zeta}_1, \tilde{\zeta}_2)$ and lemma 8 we can show that the edge inserted by **Update** satisfies property 1 (similar to the previous case).

For each triple $(\tilde{\zeta}_3, \tilde{\zeta}_4, \tilde{\zeta}_1)$ in *TCallers*, we call **Propagate**($\tilde{\zeta}_3, \tilde{\zeta}_2$). We must show that property 1 holds for $(\tilde{\zeta}_3, \tilde{\zeta}_2)$. Insertion in *TCallers* happens only at line 23, which means that $(\tilde{\zeta}_3, \tilde{\zeta}_4)$ is in *Seen*. By *IH* for $(\tilde{\zeta}_1, \tilde{\zeta}_2)$ and $(\tilde{\zeta}_3, \tilde{\zeta}_4)$ and by lemma 8, we can show that there are states $\hat{\zeta}_3$ and $\hat{\zeta}_2$ and path p' such that $\tilde{\zeta}_3 = |\hat{\zeta}_3|_{al}$, $\tilde{\zeta}_2 = |\hat{\zeta}_2|_{al}$ and $\hat{\zeta}_3 \in CE_{p'}^*(\hat{\zeta}_2)$. Hence, property 1 holds for $(\tilde{\zeta}_3, \tilde{\zeta}_2)$.

– $\tilde{\zeta}_2$ is a tail call

Similarly. □