



Context-Sensitive Demand-Driven Control-Flow Analysis

Tim Whiting^(✉) and Kimball Germane

Brigham Young University, Provo UT 84601, USA
tim@whitings.org
kimball@cs.byu.edu

Abstract. By decoupling and decomposing control flows, demand control-flow analysis (CFA) resolves only the flow segments determined necessary to produce a specified control-flow fact. It therefore presents a more flexible interface and pricing model than typical CFA, making many useful applications practical. At present, the only realization of demand CFA is the context-insensitive Demand OCFA. Typical mechanisms for adding context sensitivity are not compatible with the demand setting because the analyzer is dispatched at arbitrary program points in indeterminate contexts. We overcome this challenge by identifying a context suitable for a demand analysis and designing a representation thereof that allows it to model incomplete knowledge of the context. On top of this design, we construct Demand *m*-CFA, a context-sensitive demand CFA hierarchy. With the attractive pricing model of demand analysis and the precision offered by context sensitivity, we show that Demand *m*-CFA can replace its exhaustive counterpart in compiler backends and integrate into interactive tools such as language servers.

Keywords: Demand CFA, *m*-CFA, Context-Sensitivity, Control Flow Analysis

1 Getting into the Flow

Conventional control-flow analysis is tactless—unthinking and inconsiderate. To illustrate, consider the program fragment below which defines a recursive `fold` function. As the function iterates, it updates the index `n` and accumulator `a` using the functions `f` and `g` respectively. The values of `f` and `g` flow in parallel within `fold`, each (1) being bound in the initial call, (2) flowing to its corresponding parameter, and (3) being directly called once per iteration. But the flows of `f` and `g` don't completely overlap: `f`'s value's flow begins at `sub1` and branches into the call to `g` whereas `g`'s value's comes from a reference to `h` and only is applied within `fold`.

```
(letrec ([fold
  (λ (f g n a)
    (if (zero? n)
      a
      (fold f g (f n) (g f n a)))]])
  (fold sub1 h 42 1))
```

Now consider a tool focused on the call `(f n)` and seeking the value of `f` in order to, say, expand `f` inline. Only

the flow segments identified above respective to f are needed to fully determine this value—and know that it is fully-determined. Yet conventional control-flow analysis (CFA) is *exhaustive*, insistent on determining every segment of every flow, starting from the program’s entry point.¹ In the account it produces, the segmentation of individual flows and independence of distinct flows are completely implicit. To obtain f ’s value with a conventional CFA, the user must be willing to pay for g ’s—and any other values incidental to it—as well.

Inspired by demand dataflow analysis [4], a *demand* CFA does not determine every segment of every flow but only those segments which contribute to the values of specified program points. Moreover, because its segmentation of flows is explicit, it only need analyze each segment once and can reuse the results in any flow which contains the segment. In the `fold` example, a supporting demand CFA works backwards from the reference to f to determine its value, and considers only the three flow segments identified above to do so.

The interface and pricing model demand CFA offers make many useful applications practical. Horwitz et al. [8] identify several ways this practicality is realized:

1. One can restrict analysis to a program’s critical path or only where certain features are used.
2. One can analyze more often, and interleave analysis with other tools. For example, a demand analysis does not need to worry about optimizing transformations invalidating analysis results since one can simply re-analyze the transformed points.
3. One can let a user drive the analysis, even interactively, to enhance, e.g., an IDE experience. Having implemented our context-sensitive demand CFA in Visual Studio Code for a subset of the Koka language [10], our firsthand experience confirms its practicality here.

1.1 Adding Context Sensitivity to Demand CFA

Presently, the only realization of demand CFA is Demand 0CFA [5] which is context-insensitive. Adding context sensitivity to demand CFA would provide the same benefits that context sensitivity confers to analyses at large: increased precision and, in some cases, a reduced workload [14]. (In §2, we give an intuitive overview of Demand 0CFA and illustrate the positive effects of context sensitivity; in §4, we present a streamlined formulation of Demand 0CFA.)

However, the demand setting presents a challenge for adding context sensitivity: unlike exhaustive analyses in which the context is fully determined at each point in analysis, a demand analysis is deployed on an arbitrary program point in an undetermined context. Thus, the task of a context-sensitive demand CFA is not only to respect the context as far as it is known but also to determine

¹ Exhaustive CFA can be made to work with program components where free variables are treated specially (e.g. using Shivers’ escape technique [22, Ch. 3]). This special treatment does not change the fundamental *exhaustive* nature of the analysis nor bridge the shortcomings we describe.

unknown contexts as they are discovered relevant to analysis. We overcome this challenge through the confluence of (1) a compatible choice of context, (2) the representation of that context, and (3) that of environments, which we discuss in §5.

After addressing this challenge, we arrive at Demand *m*-CFA (§6), a hierarchy of context-sensitive demand CFA. Demand *m*-CFA determines the context only to the extent necessary to soundly answer analysis questions, as opposed to determining the entire context, which we illustrate in §2. This parsimony not only allows Demand *m*-CFA to avoid analysis work, it informs the analysis client which aspects of the context are relevant to the posed analysis question, which the client can feed into its formulation of subsequent questions.

Demand *m*-CFA is sound with respect to a concrete albeit demand semantics called *demand evaluation*, which is itself sound with respect to a standard call-by-value semantics.

We have implemented Demand *m*-CFA in several settings using the *Abstracting Definitional Interpreters* (ADI) technique [2]. We evaluate the implementation cost and performance to empirically assess Demand *m*-CFA (§8).

We conclude by discussing related (§9) and future work (§10).

2 Demand CFA, Intuitively

A user interacts with demand CFA by submitting queries, which the analyzer resolves online. There are two types of queries:

1. An *evaluation* query seeks the values to which a specified expression may evaluate. In essence, it is asking which values flow to a given expression.
2. A *trace* query seeks the sites at which the value of a specified expression may be used. In essence, it is asking where values flow from a given expression.

To resolve all but trivial queries, the analysis issues subqueries—of one type or the other—to resolve flows on which the original query depends.

We illustrate the operation of a demand CFA considering queries over the program

```
(let ([f (λ (x) x)]) (+ (f 42) (f 35)))
```

which is written in Pure Scheme (i.e. purely-functional Scheme).

2.1 Without Context Sensitivity

As many readers are likely unfamiliar with demand CFA, we'll first look at how Demand 0CFA, the context-*ins*sensitive embodiment of demand CFA, resolves queries.

Suppose that a user submits an evaluation query q_0 on the expression $(f\ 35)$. Since $(f\ 35)$ is a function application, Demand 0CFA issues a subquery q_1 to evaluate the operator f . For each procedure value of f , Demand 0CFA issues a subquery to determine the value of its body as the value of q_0 . To the left is a trace of the queries that follow q_0 .

q_0 evaluate (f 35)
 q_1 evaluate f
 q_2 evaluate (λ (x) x)
 \Rightarrow (λ (x) x)
 q_3 evaluate x
 q_4 trace (λ (x) x)
 q_5 trace f in (f 42)
 \Rightarrow (f 42)
 q_7 evaluate 42
 \Rightarrow 42
 q_6 trace f in (f 35)
 \Rightarrow (f 35)
 q_8 evaluate 35
 \Rightarrow 35

Indented queries denote subqueries whose results are used to continue resolution of the superquery. A subsequent query at the same indentation level is a query in “tail position”, whose results are those of a preceding query. A query often issues multiple queries in tail position, as this example demonstrates. The operator f is a reference, so Demand OCFA walks the syntax to find where f is bound. Upon finding it bound by a `let` expression, Demand OCFA issues a subquery q_2 to evaluate its bound expression (λ (x) x). The expression (λ (x) x) is a λ term—a value—which q_2 propagates directly to q_1 . Once q_1 receives it, Demand OCFA issues a subquery q_3 for the evaluation of its body. Its body x is a reference, so Demand OCFA walks the syntax to discover that it is λ -bound and therefore

that its value is the value of the argument at the application of (λ (x) x). That this call to (λ (x) x) originated at (f 35) is contextual information, to which Demand OCFA is insensitive. Consequently, Demand OCFA issues a trace query q_4 to find all the application sites of (λ (x) x). Because it is an expression bound to f , Demand OCFA issues a subqueries q_5 and q_6 to find the use sites of both references to f . Each of these subqueries resolves immediately since each of the references is in operator position and their results are propagated to q_4 . For each result, q_3 issues a subquery— q_7 and q_8 —to evaluate the arguments, each of which is a numeric literal, whose value is immediately known. Each query propagates its results to q_3 which propagates them to q_0 which returns them to the user. Thus, Demand OCFA concludes that (f 35) may evaluate to 42 or 35.

2.2 With Context Sensitivity

We'll now look at how Demand m -CFA, a context-sensitive demand CFA, resolves queries. As is typical, Demand m -CFA uses an environment to record the binding context of each in-scope variable. Hence, in this setting, queries and results include not only expressions, but environments as well. We will also see that Demand m -CFA does not need a timestamp to record the “current” context, a fact we discuss further in §5.5. Let's consider the same evaluation query q_0 over (f 35), this time in the top-level environment $\langle \rangle$. Like Demand OCFA, Demand m -CFA issues the subquery q_1 to determine the operator f , also in $\langle \rangle$. After it discovers (λ (x) x) to be the binding expression of f , it issues an evaluation query over it (q_2) again in $\langle \rangle$. The result of q_2 is (λ (x) x) in $\langle \rangle$, essentially

q_0 evaluate (f 35) in $\langle \rangle$
 q_1 evaluate f in $\langle \rangle$
 q_2 evaluate (λ (x) x) in $\langle \rangle$
 \Rightarrow (λ (x) x) in $\langle \rangle$
 q_3 evaluate x in \langle (f 35) \rangle
 q'_3 caller x in \langle (f 35) \rangle
 q_4 trace (λ (x) x) in $\langle \rangle$
 q_5 trace f in $\langle \rangle$
 \Rightarrow (f 42) in $\langle \rangle$
 \Rightarrow fail
 q_6 trace f in $\langle \rangle$
 \Rightarrow (f 35) in $\langle \rangle$
 \Rightarrow (f 35) in $\langle \rangle$
 q_8 evaluate 35 in $\langle \rangle$
 \Rightarrow 35 in $\langle \rangle$

a closure. As before, this result is passed first to q_1 and then to q_0 at which point Demand m -CFA constructs q_3 , an evaluation query over its body. The query's environment $\langle (\mathbf{f} \ 35) \rangle$ records the context in which the parameter \mathbf{x} was bound. In order to evaluate \mathbf{x} , Demand m -CFA issues a *caller* query q'_3 to determine the caller of $(\lambda (\mathbf{x}) \ \mathbf{x})$ that yielded the environment $\langle (\mathbf{f} \ 35) \rangle$. It then issues the trace query q_4 , this time subordinate to q'_3 , which issues q_5 and q_6 and results in the same two applications of \mathbf{f} . However, when q'_3 receives a caller from q_4 , Demand m -CFA ensures that the caller could produce the binding context of the parameter in q'_3 's environment. If so, q'_3 yields the result to q_3 ; if not, it cuts off the resolution process for that path. In this case, q_5 's result $(\mathbf{f} \ 42)$ is not compatible with q'_3 , and Demand m -CFA ceases resolving it rather than issuing q_7 . However, q_6 's result $(\mathbf{f} \ 35)$ is compatible, and its resolution continues, issuing q_8 . Resolution of q_8 occurs immediately and its result is propagated to the top-level query.

This example illustrates how Demand m -CFA uses the context information recorded in the environment to filter out discovered callers of a particular closure. Not only does this filtering increase precision in the expected way, but, in this example, it also prevents Demand m -CFA from issuing a spurious query (q_7). This behavior is an example of the well-known phenomenon of high precision keeping the analyzer's search space small [14].

2.3 ...And Indeterminacy

Each environment in the previous section was fully determined. Typically, Demand m -CFA resolves queries and produces results with environments that are—at least partially—indeterminate. For instance, to obtain all of the values

q_0 evaluate \mathbf{x} in $\langle ? \rangle$ q'_0 caller \mathbf{x} in $\langle ? \rangle$ q_1 trace $(\lambda (\mathbf{x}) \ \mathbf{x})$ in $\langle \rangle$ q_2 trace \mathbf{f} in $(\mathbf{f} \ 42)$ in $\langle \rangle$ $\Rightarrow (\mathbf{f} \ 42)$ in $\langle \rangle$ $\Rightarrow (\mathbf{f} \ 42)$ in $\langle \rangle$ q_4 evaluate \mathbf{x} in $\langle (\mathbf{f} \ 42) \rangle$ q_5 evaluate 42 in $\langle \rangle$ $\Rightarrow 42$ in $\langle \rangle$ q_3 trace \mathbf{f} in $(\mathbf{f} \ 35)$ in $\langle \rangle$ $\Rightarrow (\mathbf{f} \ 35)$ in $\langle \rangle$ q_6 evaluate \mathbf{x} in $\langle (\mathbf{f} \ 35) \rangle$ q_7 evaluate 35 in $\langle \rangle$ $\Rightarrow 35$ in $\langle \rangle$	to which \mathbf{x} , the body of $(\lambda (\mathbf{x}) \ \mathbf{x})$, may evaluate, a user may issue the query evaluate \mathbf{x} in $\langle ? \rangle$ (q_0 at left) where $?$ is a “wildcard” context to be instantiated with each context the analyzer discovers. (Though each context in the environment is indeterminate, the shape of the environment itself is determined by the lexical binding structure, which we discuss further in §5.3.) Once issued, resolution of \mathbf{x} 's evaluation again depends on a caller query q'_0 . However, because the parameter \mathbf{x} 's context is unknown, rather than filtering out callers, the caller query causes $?$ to be instantiated with a context derived from each caller. As before, Demand m -CFA dispatches a trace query q_1 which then traces occurrences of \mathbf{f} via q_2 and q_3 . This query locates the call sites $(\mathbf{f} \ 42)$ in $\langle \rangle$ and $(\mathbf{f} \ 35)$ in $\langle \rangle$. Once q_2 delivers the result $(\mathbf{f} \ 42)$ in $\langle \rangle$ to q_1 and then q'_0 , Demand m -CFA <i>instantiates</i> q_0 with this newly-discovered caller to form q_4 , whose result is q_0 's also. After creating q_3 , it continues with its resolution by
---	--

issuing q_4 to evaluate the argument 42 in $\langle \rangle$. Its result of 42 propagates from q_4 to q_0 ; The instantiation from q_3 proceeds similarly.

3 Language and Notation

We present Demand m -CFA over the unary λ calculus. It is straightforward to extend it to multiple-arity functions, data structures, constants, primitives, conditionals, and recursive forms—which we have in our implementations—but this small language suffices to discuss the essential aspects of context sensitivity.

In order to keep otherwise-identical expressions distinct, many CFA presentations uniquely label program sub-expressions.² This approach would be used, for example, to disambiguate the two references to \mathbf{f} in the program in §2. We instead use the syntactic context of each expression, which uniquely determines it, as a de-facto label, since Demand CFA consults it extensively.

In the unary λ calculus, expressions e adhere to the grammar on the left and syntactic contexts C adhere to the grammar on the right.

$$\text{Exp} \ni e ::= x \mid \lambda x.e \mid (ee) \quad \text{SynCtx} \ni C ::= \lambda x.C \mid (C e) \mid (e C) \mid \square.$$

The syntactic context C of an instance of an expression e within a program pr is pr itself with a hole \square in place of the selected instance of e . For example, the program $\lambda x.(xx)$ contains two references to x , one with syntactic context $\lambda x.(\square x)$ and the other with $\lambda x.(x \square)$.

The composition $C[e]$ of a syntactic context C with an expression e consists of C with \square replaced by e , so that $C[e]$ denotes the program with a focus on e . For example, we focus on the reference to x in operator position in the program $\lambda x.(xx)$ with $\lambda x.([x]x)$. To aid visual parsing, we also shade the context to distinguish it from the focused expression.

The composition of a context and expression is a cursor, drawn from the space $\text{Cursor} := \text{Exp} \times \text{SynCtx}$. In a cursor, we typically leave the context unspecified, referring to, e.g., a reference to x by $C[x]$ and two distinct references to x by $C_0[x]$ and $C_1[x]$ (where $C_0 \neq C_1$). The immediate syntactic context of an expression is often relevant, however, and we make it explicit by a double composition $C_0[C_1[e]]$. For example, we use $C[[x]x]$ to focus on the expression x in the operator context $(\square x)$ in the context C .

4 Demand OCFA

Demand OCFA has two modes of operation, *evaluation* and *tracing*, which users access by submitting evaluation and trace queries, respectively. An evaluation query resolves the values to which a designated expression may evaluate and a trace query resolves the sites which may apply the value of the designated

² Others operate over a form which itself names all intermediate results, such as CPS or \mathcal{A} -normal form, and identify each expression by its associated (unique) name.

expression. These modes are essentially dual and reflect the dual perspective of exhaustive CFA as an account of either (1) the λ terms $\lambda x.e$ which may be applied at each application site $(e_0 e_1)$, or (2) the application sites $(e_0 e_1)$ at which each λ term $\lambda x.e$ may be applied. However, in contrast to exhaustive CFA, Demand 0CFA resolves evaluation queries over individual expressions with a straightforward way to share work across the resolution of related queries. (It also resolves trace queries over individual expressions, but exhaustive CFAs have no counterpart to this functionality.)

The evaluation and trace modes of operation are effected by the big-step relations $\Downarrow_{eval}, \Rightarrow_{expr} \subseteq \mathbf{Cursor} \times \mathbf{Cursor}$, which are defined mutually inductively. These relations are respectively supported by the auxiliary metafunction $\mathbf{bind} \in \mathbf{Var} \times \mathbf{Cursor} \rightarrow \mathbf{Cursor}$ and relation $\Rightarrow_{find} \subseteq \mathbf{Var} \times \mathbf{Cursor} \times \mathbf{Cursor}$, which connect references to bindings and vice versa. Anticipating the addition of context sensitivity, we define \Downarrow_{eval} in terms of the relation $\Downarrow_{call} \subseteq \mathbf{Cursor} \times \mathbf{Cursor}$, which we discuss below. Figure 1 presents the definitions of all of these relations. The judg-

$$\begin{array}{c}
\Downarrow_{eval}, \Rightarrow_{expr}, \Downarrow_{call} \subseteq \mathbf{Cursor} \times \mathbf{Cursor} \qquad \Rightarrow_{find} \subseteq \mathbf{Var} \times \mathbf{Cursor} \times \mathbf{Cursor} \\
\\
\text{LAM} \qquad \qquad \qquad \text{APP} \\
\frac{C[\lambda x.e] \Downarrow_{eval} C[\lambda x.e]}{C[\lambda x.e] \Downarrow_{eval} C[\lambda x.e]} \qquad \frac{C'[(e_0 e_1)] \Downarrow_{eval} C'[\lambda x.e] \quad C'[\lambda x.[e]] \Downarrow_{eval} C_v[\lambda x.e_v]}{C'[(e_0 e_1)] \Downarrow_{eval} C_v[\lambda x.e_v]} \\
\\
\text{REF} \\
\frac{C'[e_x] = \mathbf{bind}(x, C[x]) \quad C'[e_x] \Downarrow_{call} C''[(e_0 e_1)] \quad C''[(e_0 [e_1])] \Downarrow_{eval} C_v[\lambda x.e]}{C[x] \Downarrow_{eval} C_v[\lambda x.e]} \\
\\
\text{OPERATOR} \\
\frac{}{C'[(e_0 e_1)] \Rightarrow_{expr} C'[(e_0 e_1)]} \\
\\
\text{BODY} \\
\frac{C[\lambda x.[e]] \Downarrow_{call} C'[(e_0 e_1)] \quad C'[(e_0 e_1)] \Rightarrow_{expr} C''[(e_2 e_3)]}{C[\lambda x.[e]] \Rightarrow_{expr} C''[(e_2 e_3)]} \\
\\
\text{OPERAND} \\
\frac{C'[(e_0 e_1)] \Downarrow_{eval} C'[\lambda x.e] \quad x, C'[\lambda x.[e]] \Rightarrow_{find} C_x[x] \quad C_x[x] \Rightarrow_{expr} C'[(e_2 e_3)]}{C'[(e_0 [e_1])] \Rightarrow_{expr} C'[(e_2 e_3)]} \\
\\
\text{CALL} \qquad \qquad \qquad \text{FIND-REF} \qquad \qquad \qquad \text{FIND-OPERATOR} \\
\frac{C[\lambda x.e] \Rightarrow_{expr} C'[(e_0 e_1)]}{C[\lambda x.[e]] \Downarrow_{call} C'[(e_0 e_1)]} \qquad \frac{}{x, C[x] \Rightarrow_{find} C[x]} \qquad \frac{}{x, C'[(e_0 e_1)] \Rightarrow_{find} C_x[x]} \\
\frac{}{x, C[(e_0 e_1)] \Rightarrow_{find} C_x[x]} \qquad \frac{}{x, C[(e_0 e_1)] \Rightarrow_{find} C_x[x]} \\
\\
\text{FIND-OPERAND} \qquad \qquad \qquad \text{FIND-BODY} \\
\frac{x, C[(e_0 [e_1])] \Rightarrow_{find} C_x[x]}{x, C[(e_0 e_1)] \Rightarrow_{find} C_x[x]} \qquad \frac{x \neq y \quad x, C[\lambda y.[e]] \Rightarrow_{find} C_x[x]}{x, C[\lambda y.e] \Rightarrow_{find} C_x[x]}
\end{array}$$

Fig. 1: Demand 0CFA evaluation and trace relations

ment $C[e] \Downarrow_{eval} C_v[\lambda x.e_v]$ denotes that the expression e (residing in syntactic

context C) evaluates to (a closure over) $\lambda x.e_v$. (In a context-insensitive analysis, we may represent a closure by the λ term itself.) Demand 0CFA arrives at such a judgment, as an interpreter does, by considering the type of expression being evaluated. The *Lam* rule captures the intuition that a λ term immediately evaluates to itself. The *App* rule captures the intuition that an application evaluates to whatever the body of its operator does. Hence, if the operator e_0 evaluates to $\lambda x.e$, and e evaluates to $\lambda y.e_v$, then the application $(e_0 e_1)$ evaluates to $\lambda y.e_v$ as well. Notice that the *App* does not evaluate the argument; if the argument is needed, indicated by a reference to the operator's parameter x during evaluation of its body, the *Ref* rule obtains it. The *Ref* rule captures the intuition that a reference to a parameter x takes on the value of the argument at each call site where the λ which binds x is called. The *bind* metafunction determines the binding configuration of x by walking outward on the syntax tree until it encounters x 's binder. Figure 2 presents its definition, and we note the absence of a rule for the case where x is unbound, since we define programs as closed expressions.

$\text{bind} : \text{Var} \times \text{Cursor} \rightarrow \text{Cursor}$

$$\begin{aligned} \text{bind}(x, C[[e_0] e_1]) &= \text{bind}(x, C[(e_0 e_1)]) \\ \text{bind}(x, C[e_0 [e_1]]) &= \text{bind}(x, C[(e_0 e_1)]) \\ \text{bind}(x, C[\lambda y.[e]]) &= \text{bind}(x, C[\lambda y.e]) \text{ where } y \neq x \\ \text{bind}(x, C[\lambda x.[e]]) &= C[\lambda x.[e]] \end{aligned}$$

Fig. 2: The *bind* metafunction

A judgment $C[\lambda x.[e]] \Downarrow_{\text{call}} C'[(e_0 e_1)]$ denotes that the application $(e_0 e_1)$ applies $\lambda x.e$, thereby binding x . Demand 0CFA arrives at this judgment by the *Call* rule which uses the $\Rightarrow_{\text{expr}}$ relation to determine it. In Demand 0CFA, this relation is only a thin wrapper over $\Rightarrow_{\text{expr}}$, but becomes more involved with the addition of context sensitivity.

A judgment $C[e] \Rightarrow_{\text{expr}} C'[(e_0 e_1)]$ denotes that the value of the expression e is applied at $(e_0 e_1)$. Demand 0CFA arrives at such a judgment by considering the type of the syntactic context to which the value flows. The *Operator* rule captures the intuition that, if $\lambda x.e$ flows to operator position e_0 of $(e_0 e_1)$, it is applied by $(e_0 e_1)$. The *Body* rule captures the intuition that if a value flows to the body of a λ term, then it flows to each of its callers as well. The *Operand* rule captures the intuition that a value in operand position is bound by the formal parameter of each operator value and hence to each reference to the formal parameter in the operator's body. If the operator e_f evaluates to $\lambda x.e$, then the value of e_a flows to each reference to x in e .

The $\Rightarrow_{\text{find}}$ relation associates a variable x and expression e with each reference to x in e . It's purpose is to determine where values flow locally within an expression by using and enforcing lexical scope. *Find-Ref* finds e itself if e is a reference to x . *Find-Operator* and *Find-Operand* find references to x in $(e_0 e_1)$ by searching the operator e_0 and operand e_1 , respectively. *Find-Body* finds ref-

erences to x in $\lambda x.e$ taking care that $x \neq y$ so that it does not find shadowed references.

5 Adding Context Sensitivity

A context-*insensitive* CFA is characterized by each program variable having a single entry in the store, shared by all bindings to it. In contrast, a context-sensitive CFA considers the context in which each variable is bound and ensures that only bindings made in the same context share a store entry. By extension, a context-sensitive CFA evaluates an expression under a particular environment, which identifies the context in which free variables within the expression are bound. (Like Demand 0CFA, our context-sensitive demand CFA does not materialize the store in the semantics, but it can be recovered if desired.)

A key constraint in the demand setting is that queries may be (and typically are) issued with incomplete context information. The analysis is expected both to determine the context to the extent necessary to resolve the query (but no more) and also respect the context so that the analysis is in fact context-sensitive. One implication of this expectation is that queries over expressions which are evaluated in multiple contexts should qualify results by the contexts in which they occur.

To achieve context-sensitive CFA, we must overcome three challenges: (1) identify a notion of context compatible with the demand CFA, (2) determine a representation of that context which can account for fully- or partially-indeterminate contexts, and (3) determine an appropriate representation of environments which record the context. We meet each challenge in this section.

5.1 Challenge 1: An Appropriate Notion of Context

To formulate context-sensitive demand CFA in the most general setting possible, we will avoid sensitivities to properties not present in our semantics, such as types. Since we focus on an untyped λ calculus, the most straightforward choice is call-site sensitivity.

The canonical call-site sensitivity is that of k -CFA [22] which sensitizes each binding to the last k call sites encountered in evaluation. However, this form of call-site sensitivity works against the parsimonious nature of demand CFA. To make this concrete, consider that, in the fragment `(begin (f x) (g y))`, the binding of `g`'s parameter under a last- k -call-site sensitivity depends on the particular calls made during the evaluation of `(f x)`. If the value of `(f x)` is not otherwise demanded, this dependence either provokes demand analysis to discover more of the context or requires that the portion of the context contributed by `(f x)` be left indeterminate, thereby sacrificing precision.

A better form of call-site sensitivity would allow demand CFA to discover the context through its natural course of operation. Such a form, it turns out, is the top- m -stack-frames sensitivity offered by m -CFA.

***m*-CFA’s context abstraction** The *m*-CFA hierarchy [15] is the result of an investigation into the polynomial character of *k*-CFA in an object-oriented setting versus its exponential character in a functional setting. The crucial discovery of that investigation was that object-oriented (OO) *k*-CFA induces flat environments whereas functional-oriented *k*-CFA induces nested environments. Specifically, OO *k*-CFA re-binds object (closure) variables in the allocation context which collapses the exponential environment space to a polynomial size. From this insight they derive *m*-CFA, which simulates the OO binding discipline even when analyzing functional programs; that is, when applied, *m*-CFA binds a closure’s arguments and rebinds its environment’s bindings all in the binding context of the application. Under this policy, within a given environment, all bindings are in the same context and, consequently, the analysis can represent that environment simply as that binding context.

However, this binding policy amplifies a weakness of the *k*-most-recent-call-sites abstraction of *k*-CFA. Consider a [$k = 2$]CFA analysis of a call to **f** defined by

```
(define (f x) (log "called f") (g x))
```

and suppose that **log** itself does not make any calls. When control reaches **(g x)**, the most-recent call is always **(log "called f")** so the binding context of **g**’s parameter, the most-recent two calls **(log "called f")** and **(g x)**, determine only one point in the program—the body of **f**. In *k*-CFA, only **g**’s parameter binding suffers this abbreviated precision; the bindings in **g**’s closure environment refer to the context in which they were introduced into it. In *m*-CFA, however, the bindings of **g**’s closure environment are re-bound in the context of **g**’s application as **g** is applied and once-distinct bindings are merged together in a semi-degenerate context.

To accommodate this binding policy, Might et al. [15] use the top-*m* stack frames as a binding context rather than the last *k* call sites as the former is unaffected by static sequences of calls. Hence, when control reaches **(g x)** in [$m = 2$]CFA analysis, the binding context is that of **f**’s entry and **g**’s parameter is bound in the context of **(g x)** and **f**’s caller—no context is wasted.

Because we’re using *m*-CFA’s top-*m*-stack-frames context abstraction, we call our context-sensitive demand CFA *Demand m-CFA*. It is important to keep in mind, however, that we do *not* adopt its re-binding policy. This is due to the fact that re-binding is defined in terms of a central store, which demand CFA does not explicitly model.

5.2 Challenge 2: Representing (Indeterminate) Contexts

Now that we have identified a context abstraction, we must choose a representation for it which will allow us to model incomplete knowledge. As a fully-determined context in *m*-CFA is an *m*-length sequence of call sites, a straightforward choice would be an *m*-length vector of possibly-indeterminate call sites which Demand *m*-CFA could fill in as it discovers them. However, this representation fails to capture the useful invariant that call sites within a context are

always discovered in sequence from the first (representing the top frame on the stack) to the last.

To illustrate, suppose we are evaluating x in the function $(\lambda (x) x)$ and that we have no knowledge about the binding context of x . In order to determine x 's value, we must determine the sites that call $(\lambda (x) x)$. If our analyzer determines that one such site is $(f\ 42)$, then it learns that the *top* frame—the first of the top- m frames—is $(f\ 42)$. The next frame is whatever the top frame of $(f\ 42)$'s context is, which may be indeterminate. Thus, the analyzer's knowledge of the top- m frames always increases from the top down. (This invariant holds when the analyzer enters a call as well, since it necessarily knows the call site when doing so.)

We devise a context representation that captures this invariant. A context cc^m representing m stack frames is either completely indeterminate, a pair of a call $C[(e_0\ e_1)]$ and a context cc^{m-1} of $m-1$ frames, or the stack bottom $()$ (which occurs at the top level of evaluation). A context cc^0 of zero frames (which is a different notion than the stack bottom) is simply the degenerate \square . Formally, we have

$$\text{Context}_0 \ni cc^0 ::= \square \quad \text{Context}_m \ni cc^m ::= ? \mid (C[(e_0\ e_1)], cc^{m-1}) \mid ()$$

With the context representation chosen, we can now turn to the environment representation.

5.3 Challenge 3: Representing Environments

In a lexically-scoped language, the environment at the reference x in the fragment

```
(define f (lambda (x y) ... (lambda (z) ... (lambda (w) ... x ...)) ...))
```

contains bindings for x , y , z , and w . Exhaustive CFAs typically model this environment as a finite map from variables to contexts (i.e., the type $\text{Var} \rightarrow \text{Context}$). We instead use a variable's de Bruijn index, which is statically determined by the program syntax, and model environments as a sequence Context^* which, for any finite program, has length bounded by the maximum lexical depth therein.

Given this environment representation, we make one final tweak to the definition of contexts: we will qualify an indeterminate context $?$ with the parameter of the function whose context it represents, and assume programs are α -converted—i.e. that all bound variables are unique.³ This way, even an environment of completely indeterminate contexts still determines the expression it closes. For instance, we represent the indeterminate environment of y in $(\lambda (x) ((\lambda (y) y) (\lambda (z) z)))$ by $\langle ?_y, ?_x \rangle$ which is distinct from the indeterminate environment of z , which we represent by $\langle ?_z, ?_x \rangle$, even though they have the same shape. As explained in the next section, this distinction is crucial when considering instantiating indeterminate contexts.

³ In practice, we use the syntactic context of the body instead of the parameter, which is unique even if the program is not α -converted.

5.4 Instantiating Contexts

Demand m -CFA, at certain points during resolution, discovers information about the context in which the resolved flow occurs, and must instantiate *relevant* environments with that information. For instance, in the program

```
(let ([apply (λ (f) (λ (x) (f x)))]
      (+ ((apply add1) 42)
         ((apply sub1) 35)))
```

suppose that Demand m -CFA is issued an evaluation query for $(f\ x)$ in the environment $\langle ?_x, ?_f \rangle$, i.e., the fully-indeterminate environment. With our global view, we can see that $(f\ x)$ evaluates to 43 and 34. Let's consider how Demand m -CFA would arrive at the same conclusion.

First, it would trace $(\lambda (f) (\lambda (x) (f\ x)))$ in the environment $\langle \rangle$ to the binding `apply` and then to the call sites $(\text{apply}\ \text{add1})$ in $\langle \rangle$ and $(\text{apply}\ \text{sub1})$ in $\langle \rangle$. Each of these sites provides a context in which f is bound, so the evaluation of $(f\ x)$ continues in two instantiations of $\langle ?_x, ?_f \rangle$: the first to $\langle ?_x, (\text{apply}\ \text{add1}) :: ?_{tl} \rangle$ and the second to $\langle ?_x, (\text{apply}\ \text{sub1}) :: ?_{tl} \rangle$, where $?_{tl}$ is a dummy variable for the top-level context (the stack bottom). To evaluate either `add1` or `sub1`, its argument is needed, which flows through x . Then the two callers of $(\lambda (x) (f\ x))$ must be resolved. When $(\lambda (x) (f\ x))$ flows to the result of $(\text{apply}\ \text{add1})$, it is applied immediately at $((\text{apply}\ \text{add1})\ 42)$. Similarly, when it flows to the result of $(\text{apply}\ \text{sub1})$, it is applied immediately at $((\text{apply}\ \text{sub1})\ 35)$. These two call sites provide the binding contexts for x .

We might be tempted at this point to blindly instantiate $?_x$ with each of these sites. However, in doing so, we will get *six* environments, instantiating $?_x$ in each of $\langle ?_x, ?_f \rangle$, $\langle ?_x, (\text{apply}\ \text{add1}) :: ?_{tl} \rangle$, and $\langle ?_x, (\text{apply}\ \text{sub1}) :: ?_{tl} \rangle$ with each of $((\text{apply}\ \text{add1})\ 42) :: ?_{tl}$ and $((\text{apply}\ \text{sub1})\ 35) :: ?_{tl}$. In particular, we obtain

$$\langle ((\text{apply}\ \text{sub1})\ 35) :: ?_{tl}, (\text{apply}\ \text{add1}) :: ?_{tl} \rangle \quad \langle ((\text{apply}\ \text{add1})\ 42) :: ?_{tl}, (\text{apply}\ \text{sub1}) :: ?_{tl} \rangle$$

which do not correspond to any environment which arises in an exhaustive analysis.

The issue is that blindly instantiating indeterminate contexts ignores the evaluation path taken to arrive at the call site. In particular, it ignores where the nesting λ is applied, which must be applied first (as we observed in the previous section). In this example, we must consider the context of $(\lambda (f) (\lambda (x) (f\ x)))$ before we instantiate the context of $(\lambda (x) (f\ x))$.

The solution, then, is to not substitute a indeterminate context with a more-determined context, but an entire environment headed by an indeterminate context with that same environment headed by the more-determined one. This policy would, in this example, lead to all occurrences of

$$\langle ?_x, (\text{apply}\ \text{add1}) :: ?_{tl} \rangle \quad \text{substituted with} \quad \langle ((\text{apply}\ \text{add1})\ 42) :: ?_{tl}, (\text{apply}\ \text{add1}) :: ?_{tl} \rangle$$

and

$$\langle ?_x, (\text{apply}\ \text{sub1}) :: ?_{tl} \rangle \quad \text{substituted with} \quad \langle ((\text{apply}\ \text{sub1})\ 35) :: ?_{tl}, (\text{apply}\ \text{sub1}) :: ?_{tl} \rangle$$

and no others, which is precisely what we would hope.

This policy is effective even when the result of the function does not depend on both values. For instance, when Demand m -CFA evaluates \mathbf{x} in $(\lambda (\mathbf{f}) (\lambda (\mathbf{x}) \mathbf{x}))$, it must still determine the caller of $(\lambda (\mathbf{f}) (\lambda (\mathbf{x}) \mathbf{x}))$ to determine the downstream caller of $(\lambda (\mathbf{x}) \mathbf{x})$.

Crucially, we need the indeterminate contexts to be qualified by the parameter of the function whose context it represents as mentioned previously. Without this distinction, we could instantiate both $?_x$ and $?_f$ with contexts that are only statically possible for the other. This is a clear win in terms of precision and performance, and does not distinguish more contexts than necessary due to the fact that the indeterminate variables in closure environments can be uniquely derived from their lambda.

5.5 Whence the timestamp?

In addition to introducing environments, context-sensitivity also typically introduces “timestamps” which serve as snapshots of the context at each evaluation step. For instance, classic k -CFA evaluates *configurations*, consisting of an expression, environment, and timestamp, to results.

With a top- m -stack-frames abstraction, the “current context” is simply the context of the variable(s) most recently bound in the environment. Lexical scope makes identifying these variables straightforward and our environment representation makes accessing their binding context straightforward as well. Thus, under such an abstraction, the environment uniquely determines the timestamp, and we can omit timestamps from configurations with no loss of information.

With the context, its representation, and the environment’s representation identified, we are now ready to define Demand m -CFA.

6 Demand m -CFA

$$\begin{aligned} \text{bind} &: \text{Var} \times \text{Cursor} \times \text{Env} \rightarrow \text{Cursor} \times \text{Env} \\ \text{bind}(x, C[[e_0] e_1], \hat{\rho}) &= \text{bind}(x, C[(e_0 e_1)], \hat{\rho}) \\ \text{bind}(x, C[(e_0 [e_1])], \hat{\rho}) &= \text{bind}(x, C[(e_0 e_1)], \hat{\rho}) \\ \text{bind}(x, C[\lambda y. [e]], \hat{c}c :: \hat{\rho}) &= \text{bind}(x, C[\lambda y. e], \hat{\rho}) \text{ where } y \neq x \\ \text{bind}(x, C[\lambda x. [e]], \hat{\rho}) &= (C[\lambda x. [e]], \hat{\rho}) \end{aligned}$$

Fig. 3: The bind metafunction

Demand m -CFA augments Demand 0CFA with environments and a mechanism to instantiate contexts within environments, which together provide context sensitivity. The addition of environments pervades \Downarrow_{eval}^m , \Rightarrow_{expr}^m , and \Rightarrow_{find}^m which are otherwise identical to their Demand 0CFA counterparts; these enriched relations are presented in Figure 4. The mechanism is located in \Downarrow_{call}^m , which is significantly elaborated relative to its Demand 0CFA counterpart.

$$\begin{aligned} \Downarrow_{eval}^m, \Rightarrow_{expr}^m, \Downarrow_{call}^m &\subseteq \text{Cursor} \times \text{Env} \times \text{Cursor} \times \text{Env} \\ \Rightarrow_{find}^m &\subseteq \text{Var} \times \text{Cursor} \times \text{Env} \times \text{Cursor} \times \text{Env} \end{aligned}$$

LAM

$$\frac{}{C[\lambda x.e], \hat{\rho} \Downarrow_{eval}^m C[\lambda x.e], \hat{\rho}}$$

OPERATOR

$$\frac{}{C[[e_0] e_1], \hat{\rho} \Rightarrow_{expr}^m C[(e_0 e_1)], \hat{\rho}}$$

REF

$$\frac{\begin{array}{c} (C_x[e_x], \hat{\rho}_x) = \text{bind}(x, C[x], \hat{\rho}) \\ C_x[e_x], \hat{\rho}_x \Downarrow_{call}^m C'[(e_0 e_1)], \hat{\rho}' \quad C'[(e_0 e_1)], \hat{\rho}' \Downarrow_{eval}^m C_v[\lambda x.e], \hat{\rho}_v \end{array}}{C[x], \hat{\rho} \Downarrow_{eval}^m C_v[\lambda x.e], \hat{\rho}_v}$$

APP

$$\frac{\begin{array}{c} C'[(e_0 e_1)], \hat{\rho} \Downarrow_{eval}^m C'[\lambda x.e], \hat{\rho}' \\ C'[\lambda x.[e]], \text{succ}_m(C'[(e_0 e_1)], \hat{\rho}) :: \hat{\rho}' \Downarrow_{eval}^m C_v[\lambda x.e_v], \hat{\rho}_v \end{array}}{C[(e_0 e_1)], \hat{\rho} \Downarrow_{eval}^m C_v[\lambda x.e_v], \hat{\rho}_v}$$

OPERAND

$$\frac{\begin{array}{c} C'[(e_0 e_1)], \hat{\rho} \Downarrow_{eval}^m C'[\lambda x.e], \hat{\rho}' \\ x, C'[\lambda x.[e]], \text{succ}_m(C'[(e_0 e_1)], \hat{\rho}) :: \hat{\rho}' \Rightarrow_{find}^m C_x[x], \hat{\rho}_x \\ C_x[x], \hat{\rho}_x \Rightarrow_{expr}^m C''[(e_2 e_3)], \hat{\rho}'' \end{array}}{C[(e_0 e_1)], \hat{\rho} \Rightarrow_{expr}^m C''[(e_2 e_3)], \hat{\rho}''}$$

BODY

$$\frac{\begin{array}{c} C[\lambda x.[e]], \hat{\rho} \Downarrow_{call}^m C'[(e_0 e_1)], \hat{\rho}' \quad C'[(e_0 e_1)], \hat{\rho}' \Rightarrow_{expr}^m C''[(e_2 e_3)], \hat{\rho}'' \end{array}}{C[\lambda x.[e]], \hat{\rho} \Rightarrow_{expr}^m C''[(e_2 e_3)], \hat{\rho}''}$$

FIND-REF

$$\frac{}{x, C[x], \hat{\rho} \Rightarrow_{find}^m C[x], \hat{\rho}}$$

FIND-OPERATOR

$$\frac{}{\begin{array}{c} x, C'[(e_0 e_1)], \hat{\rho} \Rightarrow_{find}^m C_x[x], \hat{\rho}_x \\ x, C[(e_0 e_1)], \hat{\rho} \Rightarrow_{find}^m C_x[x], \hat{\rho}_x \end{array}}$$

FIND-OPERAND

$$\frac{}{\begin{array}{c} x, C'[(e_0 e_1)], \hat{\rho} \Rightarrow_{find}^m C_x[x], \hat{\rho}_x \\ x, C[(e_0 e_1)], \hat{\rho} \Rightarrow_{find}^m C_x[x], \hat{\rho}_x \end{array}}$$

FIND-BODY

$$\frac{}{\begin{array}{c} x \neq y \quad x, C[\lambda y.[e]], ?_y :: \hat{\rho} \Rightarrow_{find}^m C_x[x], \hat{\rho}_x \\ x, C[\lambda y.e], \hat{\rho} \Rightarrow_{find}^m C_x[x], \hat{\rho}_x \end{array}}$$

Fig. 4: Demand m -CFA Resolution

In Demand m -CFA, environments are constructed when the analysis follows evaluation forward, such as entering a call, and instantiated when the analysis follows it backward, such as when the caller of an entry configuration is sought.

6.1 Following Evaluation Forwards

When a call is entered, which occurs in the *App* and *Operand* rules, a new environment is synthesized using the succ_m metafunction which determines the binding context of the call as

$$\text{succ}_m(\mathcal{C}[(e_0 e_1)], \hat{c}c :: \hat{\rho}) = \mathcal{C}[(e_0 e_1)] :: \hat{c}c]_m$$

where $[\cdot]_m$ is defined

$$[\hat{c}c]_0 = \square \quad [?_x]_m = ?_x \quad \mathcal{C}[(e_0 e_1)] :: \hat{c}c]_m = \mathcal{C}[(e_0 e_1)] :: [\hat{c}c]_{m-1}$$

6.2 Following Evaluation Backwards

When the value of a reference is demanded, Demand m -CFA first uses the *bind* metafunction to locate its binding configuration. Its definition, presented in Figure 3, is lifted from Demand 0CFA's to accommodate environments.

With the binding configuration in hand, Demand m -CFA issues a call query to resolve calls which enter that configuration. The \Downarrow_{call}^m relation resolves such queries. As in Demand 0CFA, \Downarrow_{call}^m defers to \Rightarrow_{expr}^m to determine the call configurations at which a particular closure is applied. In the presence of contexts, however, there are two possibilities when a call configuration is resolved. The first possibility is that the call entry context computed from that configuration precisely matches that of the binding configuration, in which case the call configuration is a caller of the binding configuration; we refer to this as the *Known-Call* case. The second possibility is that the call entry context refines the binding configuration's context, in which case the refined context and environment should be instantiated to the refining context and environment; we refer to this as the *Unknown-Call* case.

6.3 Discovering Callers

We now describe how \Downarrow_{call}^m , presented in Figure 5, handles each of these cases. Each case is handled by a corresponding rule. Each rule is predicated on the *reachability* of the call query, which we discuss shortly, and resolution of the corresponding trace query. The *Known-Call* rule says that, if the call entry context of the delivered call matches the binding configuration's, then the delivered call is a known call—*known* because the caller query has the context of the call already in its environment. If $\hat{c}c_1 \neq \hat{c}c_0$, however, then the result constitutes an *unknown* caller. In this case, *Unknown-Call* considers whether $\hat{c}c_1$ refines $\hat{c}c_0$

$$\begin{array}{c}
\text{KNOWN-CALL} \\
\frac{C[\lambda x.e], \hat{\rho} \Rightarrow_{\text{expr}}^m C'[(e_0 e_1)], \hat{\rho}' \quad q \uparrow_{\text{reach}}^m \text{call}(C[\lambda x.[e]], \hat{c}c_0 :: \hat{\rho}) \quad \hat{c}c_1 := \text{succ}_m(C'[(e_0 e_1)], \hat{\rho}') \quad \hat{c}c_1 = \hat{c}c_0}{C[\lambda x.[e]], \hat{c}c_0 :: \hat{\rho} \Downarrow_{\text{call}}^m C'[(e_0 e_1)], \hat{\rho}'} \\
\\
\text{UNKNOWN-CALL} \\
\frac{C[\lambda x.e], \hat{\rho} \Rightarrow_{\text{expr}}^m C'[(e_0 e_1)], \hat{\rho}' \quad q \uparrow_{\text{reach}}^m \text{call}(C[\lambda x.[e]], \hat{c}c_0 :: \hat{\rho}) \quad \hat{c}c_1 := \text{succ}_m(C'[(e_0 e_1)], \hat{\rho}') \quad \hat{c}c_1 \sqsubset \hat{c}c_0}{\hat{c}c_1 :: \hat{\rho} / \hat{c}c_0 :: \hat{\rho}}
\end{array}$$

Fig. 5: Demand m -CFA Call Discovery

in the sense that $\hat{c}c_0$ can be instantiated to form $\hat{c}c_1$. Formally, the refinement relation \sqsubset is defined as the least relation satisfying

$$C'[(e_0 e_1)] :: \hat{c}c \sqsubset ?_{C[e]} C'[(e_0 e_1)] :: \hat{c}c_1 \sqsubset C'[(e_0 e_1)] :: \hat{c}c_0 \iff \hat{c}c_1 \sqsubset \hat{c}c_0$$

If $\hat{c}c_1$ refines $\hat{c}c_0$, *Unknown-Call* does not conclude a $\Downarrow_{\text{call}}^m$ judgment, but rather an *instantiation* judgment $\hat{c}c_1 :: \hat{\rho} / \hat{c}c_0 :: \hat{\rho}$ which denotes that *any* environment $\hat{c}c_0 :: \hat{\rho}$ may be instantiated to $\hat{c}c_1 :: \hat{\rho}$. It is by this instantiation that *Known-Call* will be triggered. When $\hat{c}c_1$ does not refine $\hat{c}c_0$, the resultant caller is ignored which, in effect, filters the callers to only those which are compatible and ensures that Demand m -CFA is indeed context-sensitive.

The $\Downarrow_{\text{call}}^m$ relation relies on a reachability relation $\uparrow_{\text{reach}}^m$ which establishes which queries are (transitively) issued in the course of resolving a top-level query. This relation allows $\Downarrow_{\text{call}}^m$ to instantiate only seen queries with refinement results. This relation is defined over queries themselves; the judgment $q \uparrow_{\text{reach}}^m q'$ captures that, if query q is reachable in analysis, then q' is also, where q is of the form

$$\text{Query} \ni q ::= \text{eval}(C[e], \hat{\rho}) \mid \text{expr}(C[e], \hat{\rho}) \mid \text{call}(C[e], \hat{\rho})$$

Figure 6 presents a formal definition of $\uparrow_{\text{reach}}^m$. The *Reflexivity* rule ensures that the top-level query is considered reachable. The *Ref-Caller* and *Ref-Argument* rules establish reachability corresponding to the *Ref* rule of $\Downarrow_{\text{eval}}^m$: *Ref-Caller* makes the caller query reachable and, if it succeeds, *Ref-Argument* makes the ensuing evaluation query reachable. *App-Operator* and *App-Body* do the same for the *App* rule of $\Downarrow_{\text{eval}}^m$, making, respectively, the operator evaluation query reachable and, if it yields a value, the body evaluation query reachable. *Operand-Operator* makes the evaluation query of the *Operand* rule reachable and *Operand-Body* makes the trace query of any references in the operator body reachable. *Body-Caller-Find* makes the caller query of *Body* reachable; if a caller is found, *Body-Caller-Trace* makes the trace query of that caller reachable. Finally, *Call-Trace* makes sure that the trace query of an enclosing λ of a caller query is reachable.

Figure 7 presents an extension of $\uparrow_{\text{reach}}^m$ which propagates instantiations to evaluation and trace queries. The *Instantiate-Reachable-** rules ensure that if a

$$\begin{array}{c}
\text{REFLEXIVITY} \\
\frac{}{q \uparrow_{\text{reach}}^m q} \\
\\
\text{REF-ARGUMENT} \\
\frac{q \uparrow_{\text{reach}}^m \text{eval}(C[x], \hat{\rho}) \quad (C'[e], \hat{\rho}') = \text{bind}(x, C[x], \hat{\rho})}{q \uparrow_{\text{reach}}^m \text{eval}(C''[(e_0 \ e_1)], \hat{\rho}'')} \\
\\
\text{REF-CALLER} \\
\frac{q \uparrow_{\text{reach}}^m \text{eval}(C[x], \hat{\rho}) \quad (C'[e], \hat{\rho}') = \text{bind}(x, C[x], \hat{\rho})}{q \uparrow_{\text{reach}}^m \text{call}(C'[e], \hat{\rho}')} \\
\\
\text{APP-OPERATOR} \\
\frac{q \uparrow_{\text{reach}}^m \text{eval}(C[(e_0 \ e_1)], \hat{\rho})}{q \uparrow_{\text{reach}}^m \text{eval}(C[[e_0] \ e_1], \hat{\rho})} \\
\\
\text{APP-BODY} \\
\frac{q \uparrow_{\text{reach}}^m \text{eval}(C[(e_0 \ e_1)], \hat{\rho}) \quad C[[e_0] \ e_1], \hat{\rho} \Downarrow_{\text{eval}}^m C'[\lambda x. e], \hat{\rho}'}{q \uparrow_{\text{reach}}^m \text{eval}(C'[\lambda x. [e]], \text{succ}_m(C[(e_0 \ e_1)], \hat{\rho}) :: \hat{\rho}')} \\
\\
\text{CALL-TRACE} \\
\frac{q \uparrow_{\text{reach}}^m \text{call}(C[\lambda x. [e]], \hat{c}c :: \hat{\rho})}{q \uparrow_{\text{reach}}^m \text{expr}(C[\lambda x. e], \hat{\rho})} \\
\\
\text{OPERAND-OPERATOR} \\
\frac{q \uparrow_{\text{reach}}^m \text{expr}(C[(e_0 \ e_1)], \hat{\rho})}{q \uparrow_{\text{reach}}^m \text{eval}(C[[e_0] \ e_1], \hat{\rho})} \\
\\
\text{OPERAND-BODY} \\
\frac{q \uparrow_{\text{reach}}^m \text{expr}(C[(e_0 \ e_1)], \hat{\rho}) \quad C[[e_0] \ e_1], \hat{\rho} \Downarrow_{\text{eval}}^m C'[\lambda x. e], \hat{\rho}' \quad x, C'[\lambda x. [e]], \text{succ}_m(C[(e_0 \ e_1)], \hat{\rho}) :: \hat{\rho}' \Rightarrow_{\text{find}}^m C_x[x], \hat{\rho}_x}{q \uparrow_{\text{reach}}^m \text{expr}(C_x[x], \hat{\rho}_x)} \\
\\
\text{BODY-CALLER-FIND} \\
\frac{q \uparrow_{\text{reach}}^m \text{expr}(C[\lambda x. [e]], \hat{\rho})}{q \uparrow_{\text{reach}}^m \text{call}(C[\lambda x. [e]], \hat{\rho})} \\
\\
\text{BODY-CALLER-TRACE} \\
\frac{q \uparrow_{\text{reach}}^m \text{expr}(C[\lambda x. [e]], \hat{\rho}) \quad C[\lambda x. [e]], \hat{\rho} \Downarrow_{\text{call}}^m C'[(e_0 \ e_1)], \hat{\rho}'}{q \uparrow_{\text{reach}}^m \text{expr}(C'[(e_0 \ e_1)], \hat{\rho}')}
\end{array}$$

Fig. 6: Demand m -CFA Reachability

$$\begin{array}{c}
\text{INSTANTIATE-REACHABLE-EVAL} \\
\frac{q \uparrow_{reach}^m \text{eval}(C[e], \hat{\rho}) \quad \hat{\rho}_1 / \hat{\rho}_0}{q \uparrow_{reach}^m \text{eval}(C[e], \hat{\rho}[\hat{\rho}_1 / \hat{\rho}_0])}
\end{array}
\qquad
\begin{array}{c}
\text{INSTANTIATE-EVAL} \\
\frac{\hat{\rho}_1 / \hat{\rho}_0 \quad C[e], \hat{\rho}[\hat{\rho}_1 / \hat{\rho}_0] \Downarrow_{eval}^m C_v[\lambda x.e], \hat{\rho}_v}{C[e], \hat{\rho} \Downarrow_{eval}^m C_v[\lambda x.e], \hat{\rho}_v}
\end{array}$$

$$\begin{array}{c}
\text{INSTANTIATE-REACHABLE-EXPR} \\
\frac{q \uparrow_{reach}^m \text{expr}(C[e], \hat{\rho}) \quad \hat{\rho}_1 / \hat{\rho}_0}{q \uparrow_{reach}^m \text{expr}(C[e], \hat{\rho}[\hat{\rho}_1 / \hat{\rho}_0])}
\end{array}
\qquad
\begin{array}{c}
\text{INSTANTIATE-EXPR} \\
\frac{\hat{\rho}_1 / \hat{\rho}_0 \quad C[e], \hat{\rho}[\hat{\rho}_1 / \hat{\rho}_0] \Rightarrow_{expr}^m C'[(e_0 e_1)], \hat{\rho}'}{C[e], \hat{\rho} \Rightarrow_{expr}^m C'[(e_0 e_1)], \hat{\rho}'}
\end{array}$$

$$\begin{array}{c}
\text{INSTANTIATE-REACHABLE-CALL} \\
\frac{q \uparrow_{reach}^m \text{call}(C[e], \hat{\rho}) \quad \hat{\rho}_1 / \hat{\rho}_0}{q \uparrow_{reach}^m \text{call}(C[e], \hat{\rho}[\hat{\rho}_1 / \hat{\rho}_0])}
\end{array}
\qquad
\begin{array}{c}
\text{INSTANTIATE-CALL} \\
\frac{\hat{\rho}_1 / \hat{\rho}_0 \quad C[e], \hat{\rho}[\hat{\rho}_1 / \hat{\rho}_0] \Downarrow_{call}^m C'[(e_0 e_1)], \hat{\rho}'}{C[e], \hat{\rho} \Downarrow_{call}^m C'[(e_0 e_1)], \hat{\rho}'}
\end{array}$$

$$\begin{array}{c}
\text{APP-BODY-INSTANTIATION} \\
\frac{q \uparrow_{reach}^m \text{eval}(C[(e_0 e_1)], \hat{\rho}) \quad C[(e_0 e_1)], \hat{\rho} \Downarrow_{eval}^m C'[\lambda x.e], \hat{\rho}'}{\text{succ}_m(C[(e_0 e_1)], \hat{\rho}) :: \hat{\rho}' / ?_x :: \hat{\rho}'}
\end{array}$$

$$\begin{array}{c}
\text{OPERAND-BODY-INSTANTIATION} \\
\frac{q \uparrow_{reach}^m \text{expr}(C[(e_0 [e_1])], \hat{\rho}) \quad C[(e_0 [e_1])], \hat{\rho} \Downarrow_{eval}^m C'[\lambda x.e], \hat{\rho}'}{\text{succ}_m(C[(e_0 e_1)], \hat{\rho}) :: \hat{\rho}' / ?_x :: \hat{\rho}'}
\end{array}$$

Fig. 7: Demand m -CFA Instantiation

query of any kind is reachable, then its instantiation is too. When an instantiation $\hat{\rho}_1 / \hat{\rho}_0$ does not apply (so that $\hat{\rho}$ is unchanged), each rule reduces to a trivial inference. The counterpart *Instantiate-** rules, also present in Figure 7, each extend one of \Downarrow_{eval}^m , \Rightarrow_{expr}^m , and \Downarrow_{call}^m so that, if an instantiated query of that type yields a result, the original, uninstantiated query yields that same result. As discussed at the beginning of this section, Demand m -CFA also discovers instantiations when it extends the environment in the *App* and *Operand* rules. The *App-Body-Instantiation* and *Operand-Body-Instantiation* rules capture these cases.

The definition of Demand m -CFA in terms of an “evaluation” relation (which includes evaluation, trace, and caller resolution) and a reachability relation follows the full formal approach of *Abstracting Definitional Interpreters* by Darais [3]. From this correspondence, we can define the Demand m -CFA resolution of a given query as the least fixed point of these relations, effectively computable with the algorithm Darais [3] provides. We discuss this implementation in more depth in §8.

7 Demand m -CFA Correctness

Demand m -CFA is a hierarchy of demand CFA. Instances higher in the hierarchy naturally have larger state spaces. The size $|\Downarrow_{eval}^m|$ of the \Downarrow_{eval}^m relation satisfies

the inequality

$$|\Downarrow_{eval}^m| \leq |Config \times Config| = |Config|^2 = |Exp \times Env|^2 = |Exp|^2 |Env|^2 = n^2 |Env|^2$$

where n is the size of the program. We then have

$$|Env| \leq |Ctx|^n \leq (|Call| + 1)^{mn} \leq n^{mn}$$

since the size of environments is statically bound and may be indeterminate. Thus, $|\Downarrow_{eval}^m| \leq n^{mn+2}$; the state space of \Downarrow_{eval}^m is finite but with an exponential bound; the state spaces of \Rightarrow_{expr}^m and \Downarrow_{call}^m behave similarly.

7.1 Demand m -CFA Refinement

Instances higher in the hierarchy are also more precise, which we formally express with the following theorems.

Theorem 1 (Evaluation Refinement). *If $C[e], \hat{\rho}_0 \Downarrow_{eval}^{m+1} C_v[\lambda x.e_v], \hat{\rho}'_0$ where $\hat{\rho}_0 \sqsubseteq \hat{\rho}_1$ then $C[e], \hat{\rho}_1 \Downarrow_{eval}^m C_v[\lambda x.e_v], \hat{\rho}'_1$ where $\hat{\rho}'_0 \sqsubseteq \hat{\rho}'_1$.*

Theorem 2 (Trace Refinement). *If $C[e], \hat{\rho}_0 \Rightarrow_{expr}^{m+1} C'[(e_0 e_1)], \hat{\rho}'_0$ where $\hat{\rho}_0 \sqsubseteq \hat{\rho}_1$ then $C[e], \hat{\rho}_1 \Rightarrow_{expr}^m C'[(e_0 e_1)], \hat{\rho}'_1$ where $\hat{\rho}'_0 \sqsubseteq \hat{\rho}'_1$.*

Theorem 3 (Caller Refinement). *If $C[e], \hat{\rho}_0 \Downarrow_{call}^{m+1} C'[(e_0 e_1)], \hat{\rho}'_0$ where $\hat{\rho}_0 \sqsubseteq \hat{\rho}_1$ then $C[e], \hat{\rho}_1 \Downarrow_{call}^m C'[(e_0 e_1)], \hat{\rho}'_1$ where $\hat{\rho}'_0 \sqsubseteq \hat{\rho}'_1$.*

These theorems state that refining configurations submitted to Demand m -CFA and its successor Demand $m + 1$ -CFA yield refining configurations. The proof proceeds directly (if laboriously) by induction on the derivations of the relations.

7.2 Demand ∞ -CFA and Demand Evaluation

To show that Demand m -CFA is sound with respect to a standard call-by-value (CBV) semantics, we consider the limit of the hierarchy, Demand ∞ -CFA, in which context lengths are unbounded. From here, we bridge Demand ∞ -CFA to a CBV semantics with a concrete form of demand analysis called *Demand Evaluation*. Our strategy is to show that the Demand ∞ -CFA semantics is equivalent to Demand Evaluation which itself is sound with respect to a standard CBV semantics. Demand Evaluation environments, rather than being a sequence of contexts, are sequences of addresses. Like contexts, an address may denote an indeterminate context (i.e. call) which manifests as an address which is not mapped in the store. A store is a pair consisting of a map from addresses to calls and the next address to use; the initial store is $(\perp, 0)$.

Now it is straightforward to express the equivalence between the Demand ∞ -CFA relations and Demand Evaluation.

Theorem 4 (Evaluation Equivalence). *If $\hat{\rho}_0 \hat{\rho} \Leftrightarrow_{\rho} \rho_0, \sigma_0$ then $C[e], \hat{\rho}_0 \Downarrow_{eval}^{\infty} C'[\lambda x.e], \hat{\rho}_1$ iff $C[e], \rho_0, \sigma_0 \Downarrow_{eval}^{de} C'[\lambda x.e], \rho_1, \sigma_1$ where $\hat{\rho}_1 \hat{\rho} \Leftrightarrow_{\rho} \rho_1, \sigma_1$.*

Theorem 5 (Trace Equivalence). *If $\hat{\rho}_0 \hat{\rho} \Leftrightarrow_{\rho} \rho_0, \sigma_0$ then $C[e], \hat{\rho}_0 \Rightarrow_{\text{expr}}^{\infty} C'[(e_0 e_1)], \hat{\rho}_1$ iff $C[e], \rho_0, \sigma_0 \Rightarrow_{\text{expr}}^{\text{de}} C'[(e_0 e_1)], \rho_1, \sigma_1$ where $\hat{\rho}_1 \hat{\rho} \Leftrightarrow_{\rho} \rho_1, \sigma_1$.*

Theorem 6 (Caller Equivalence). *If $\hat{\rho}_0 \hat{\rho} \Leftrightarrow_{\rho} \rho_0, \sigma_0$ then $C[e], \hat{\rho}_0 \Downarrow_{\text{call}}^{\infty} C'[(e_0 e_1)], \hat{\rho}_1$ iff $C[e], \rho_0, \sigma_0 \Downarrow_{\text{call}} C'[(e_0 e_1)], \rho_1, \sigma_1$ where $\hat{\rho}_1 \hat{\rho} \Leftrightarrow_{\rho} \rho_1, \sigma_1$.*

These theorems are proved by induction on the derivations, corresponding instantiation of environments on the Demand ∞ -CFA side with mapping an address on the Demand Evaluation side. The technical report [30] presents Demand ∞ -CFA and demand evaluation formally.

8 Evaluation

We implemented Demand m -CFA for a subset of R6RS Scheme [24] including `let`, `let*`, `letrec` binding forms, mutually-recursive definitions and a few dozen primitives. We also implemented support for constructors, numbers, symbols, strings, and characters. We call this language Pure Scheme. Although our analysis does not handle imperative constructs, the analysis can still be deployed in programs with imperative features. It produces sound results for flows that do not experience imperative update and soundly detect when flows do.

We used the *Abstracting Definitional Interpreters* approach [2](ADI) to implement m -CFA and Demand m -CFA analyses for Pure Scheme.

The results of the analysis are obtained through ADI's memoized fixpoint cache. This cache gives control flow results for each segment of control flow independently keyed by the query or subquery. Additionally the cache also has a key for each indeterminate environment which maps to each discovered environment instantiation reachable from the top level query.

We evaluate Demand m -CFA with respect to the following questions:

1. How many queries can Demand m -CFA resolve at a given level of precision (m) and effort allocated per query?
2. How many of those queries yield answers precise enough to support inlining?
3. How does the implementation cost compare to a typical CFA?

To answer the first two questions, we evaluate Demand m -CFA on a set of R6RS programs used by Johnson et al. [9], which is standard within the literature.

8.1 Query Resolution Rate

Like an exhaustive analysis, a demand analysis is subject to client-imposed resource constraints. However, unlike with exhaustive analysis, failure in a demand analysis is not catastrophic: because analysis information is sought selectively and independently through queries, the client can increase the budget and reissue the query or proceed without control flow information.

We evaluate the proportion of queries that Demand m -CFA is able to resolve within specific effort limits and at different context sensitivity levels. Effort is measured in gas, where each subquery consumes one unit.

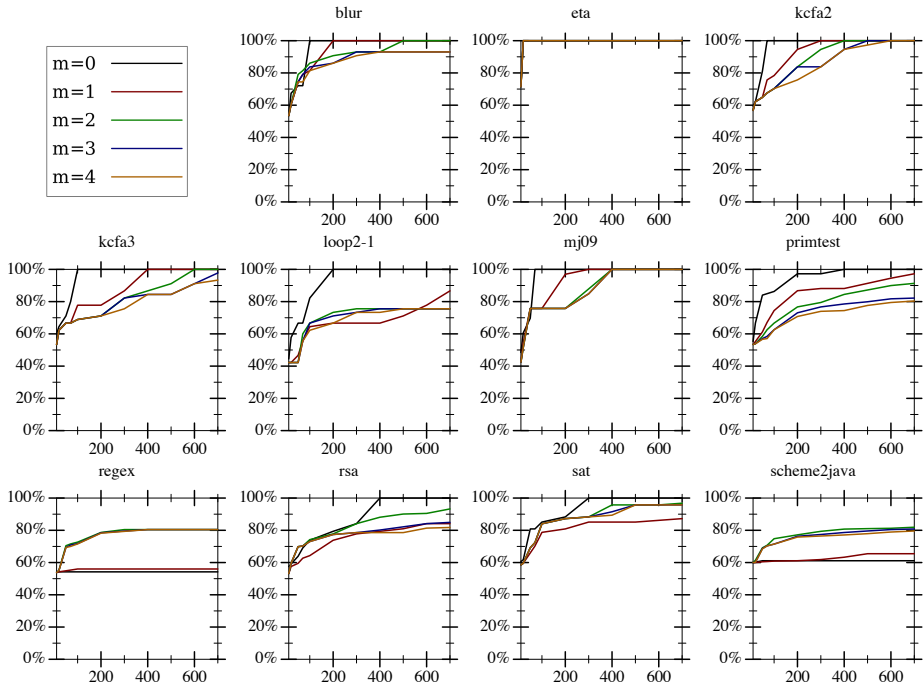


Fig. 8: The percent of queries answered (y-axis) given the per-query gas allocated to Demand m -CFA (x-axis). Many trivial queries (40%-60%) are answered given very little effort, as seen by the immediate jump. Some additional queries are answered with more effort, but some non-trivial queries are answered with less effort by increasing the context sensitivity parameter m (as seen in the larger `regex` and `scheme2java` example programs). Increasing context sensitivity can sometime help the analysis distinguish context-irrelevant flows and avoid spending gas on unrelated parts of the program.

Figure 8 displays the percent of queries answered depending on the effort allotted. The graphs trend upward and to the right, regardless of our choice of m . Notable exceptions include `regex`, and `scheme2java` which both contain highly connected control flow graphs. Both significantly improve when explored at higher m which exemplifies the known paradox of precision from exhaustive CFA, in which increasing precision can sometimes decrease the state space due to fewer spurious flows being merged. Our results show that this phenomenon holds even for subportions of programs. We hypothesize that this paradox might be more useful in the demand context because Demand m -CFA (1) does not attempt to explore undemanded (and therefore query-irrelevant) parts of the program at high levels of context-sensitivity and cost, and (2) only explores the environment instantiations that are required to answer the query, leaving as much of the context indeterminate as possible.

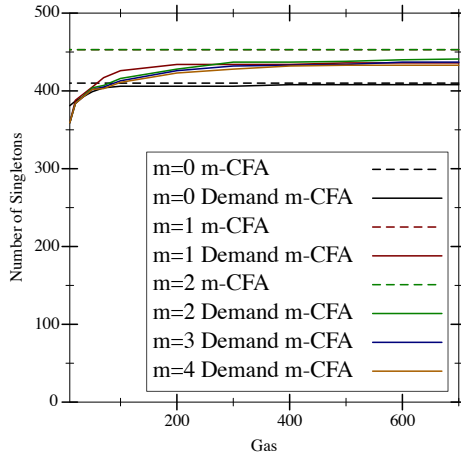


Fig. 9: The number of singleton flow sets (y-axis) found by a Demand m -CFA analysis given the gas allocated per query (x-axis) aggregated across programs. Dashed lines represent the baseline number of singleton flow sets found by an exhaustive exponential m -CFA analysis (graphed as a horizontal line). Dashed lines above $m = 2$ are not shown because `scheme2java` timed out for $m \geq 2$ exhaustive m -CFA. The dashed line for $m = 1$ is underneath $m = 2$.

8.2 Actionable Query Resolution Rate

For many clients, query results must meet a certain level of precision to be useful. For instance, compilers require results to be singleton flow sets in order to inline or otherwise optimize programs. Figure 9 shows the number of singleton flow sets Demand m -CFA was able to obtain for given levels of precision (m) and effort (gas) across the entire program corpus. The total number of singleton flow sets obtained by the corresponding exhaustive m -CFA for each level of precision is also graphed. (Because one program exceeded a 24-hour timeout we do not include results of exhaustive m -CFA for $m > 2$. A breakdown of results per program is given in the technical report [30].)

The data show that Demand m -CFA generally finds most of the singleton value flows that exhaustive m -CFA does, and support that most actionable control flow information is available at low effort bounds. The data also indicate that the attained level of context sensitivity is mostly independent of the effort allotted. This is consistent with Demand m -CFA's ability to avoid instantiating indeterminate aspects of environments when they are not relevant. In contrast, exhaustive m -CFA fully instantiates all contexts which often leads to an explosion in the explored state space.

As Figure 9 shows, Demand m -CFA is not able to discover all the singleton flow sets as the exhaustive m -CFA at any level of effort bound tested. Because of its reachability assumptions, Demand m -CFA finds more flows than exhaustive

m-CFA, which reduces the number of singleton flows it detects. For instance, if Demand *m*-CFA is finding the callers of `f` in the expression `(λ (g) (f 42))` to discover arguments to `f`, it assumes that `(f 42)` is reachable—i.e., it assumes that `(λ (g) (f 42))` is called. If that assumption is false, then the argument `42` does not actually contribute to the value that Demand *m*-CFA is resolving, and its inclusion is manifest as imprecision. The data indicate, that in practice, this reduction is less than 5% across the program corpus.

8.3 Implementation Cost

There are two notions of cost that we examine in this section: (1) the raw cost of implementation on a small core language relative to a corresponding exhaustive analysis, and (2) the actual cost of implementation for a real-world language.

To evaluate the first notion of cost, we implemented both exhaustive and Demand *m*-CFA for Pure Scheme. The codebases shared approximately 1000 lines of code. Exhaustive *m*-CFA required approximately 450 lines of additional code. Demand *m*-CFA required approximately 630 lines of additional code. For Pure Scheme the total cost of implementing Demand *m*-CFA is not significantly higher than that of exhaustive *m*-CFA.

To evaluate the second notion of cost, we implemented Demand *m*-CFA for the subset of the Koka language which does not include algebraic effect handlers and mutation. This highlights a strength of Demand *m*-CFA that, despite being implemented for only a subset of Koka, our implementation is still useful. For example, we integrated it into a language server which is able to provide low-latency control flow information to clients. In contrast, to implement an exhaustive analysis for the same purpose, one must implement the full set of Koka’s features. While exhaustive analyses support mutation without issue, support for algebraic effect handlers is not currently tractable. Our implementation, written in Haskell, consists of approximately 3000 lines of code.

Our experience with the Koka compiler indicates that integrating Demand *m*-CFA into a real-world compiler or language server is in some cases *more* tractable than an exhaustive analysis, since not all language features or primitives need to be supported before getting useful and actionable results.⁴

9 Related Work

The original inspiration for demand CFA is demand dataflow analysis [8] which refers to dataflow analysis algorithms which allow selective, local, parsimonious analysis of first-order programs driven by the user. Demand CFA refers to a class of algorithms with those same characteristics which operate in the presence of first-class functions. This work extends Demand OCFA [5], currently the sole embodiment of demand CFA, with context sensitivity using a context abstraction similar to that of *m*-CFA [15].

⁴ The Pure Scheme and Koka implementations are available on GitHub[29,28].

The most-closely related work is that of Schoepe et al. [20] which utilizes first-order forward and backward analyses to construct call graphs on demand to compose a higher-order analysis. Although it treats control with context sensitivity, it does not model or track the environment structure of higher-order objects. Our work has a similar goal—to achieve demand-driven analysis of higher-order programs—but it has sought to provide a concise analysis which models both the control and environment structure of evaluation with equal facility.

DDPA [17] is nominally a context-sensitive, demand-driven analysis for higher-order programs. However, before resolving any on-demand queries, DDPA must bootstrap a global control-flow graph to support them. Because of this large, fixed, up-front cost, DDPA does not provide the pricing model expected of a demand analysis.

Several other “demand-driven” analyses exist for functional programs. Midtgaard and Jensen [13] present a “demand-driven 0-CFA” derived via a calculational approach which analyzes only those parts of the program on which the *program’s* result depends. Biswas [1] presents a demand analysis that takes a similar approach to “demand-driven 0-CFA” to analyze *first-order* functional programs. These analyses are certainly “demand-driven” in the sense the authors intend (demanded by the full program), but not in the sense that we use it, as a descriptor for generic demand CFA. Our Demand CFA can be considered a loose extension in spirit which analyzes only those parts of a *higher-order* functional program on which *a selected expression’s* result depends.

Nicolay et al. [16] produce a flow analysis where new information is propagated between interprocedural and intraprocedural parts of the program via the store. As such they incrementally build a reactive control flow graph on demand. In addition to read and write dependency tracking their analysis explicitly tracks call effects which triggers a demand for functions stored at that point in the store to be evaluated interprocedurally with new arguments. However, while they present the theory for a context sensitive analysis, they do not present results for context sensitivity, and while their approach could be adapted to a demand based setting, their focus is on full program analysis.

Heintze and McAllester [6] describe the “demand-driven” subtransitive CFA which computes an underapproximation of control flow in linear time and can be transitively closed for an additional quadratic cost. Their analysis exploits type structure and applies to typed programs with bounded type, whereas our formulation neither considers nor requires types.

Points-to analysis is the analogue of CFA in an object-oriented setting in the sense that both are fundamental analyses that provide the necessary support for higher-level analysis. Many context-sensitive demand-driven points-to analyses (e.g. Lu et al. [12] Shang et al. [21] Späth et al. [23] Su et al. [26]) exist, formulated for Java. Though both points-to analysis and CFA target higher-order programs, Might et al. [15] observed that the explicit object creation in the object-oriented setting induces flat closures whereas implicit closure creation in the functional setting induces nested closures. Moreover, mutation is routine in many object-oriented settings (e.g. Java) in which points-to analyses are ex-

pressed. Both nested environments and the prevalence of immutable variables are fundamental to our realization of demand CFA.

Pointer analysis is the imperative analogue of control-flow analysis and it too enjoys a rich literature. When function pointers are present, a demand-driven pointer analysis must be able to reconstruct the call graph on the fly, a requirement shared by demand-driven CFA.

Heintze and Tardieu [7] present a demand-driven pointer analysis for C capable of constructing the call graph on the fly. They recognize that most call targets are not specified indirectly through pointers and advocate demand-driven analysis to resolve indirect targets when they appear. In the core language of Demand *m*-CFA, calls with indirect targets are commonplace; the language was chosen to emphasize those cases in particular. In a full-featured language, calls with indirect targets—though still more common than in C—are less common, since calls to primitives are typically direct and statically discernible. However, closures are not merely code pointers, but environments too, and we have had to exercise great care to ensure that Demand *m*-CFA correctly models their behavior.

Saha and Ramakrishnan [19] formulate points-to analysis of C as a logic program, so that the incremental-evaluation capabilities of the logic engine yield an incremental analysis. They combine an incremental and demand-driven approach to update the points-to model in response to changes in the program. Our work is similar in that it may be possible to cast it as a logic program and thereby combine an initial exhaustive CFA with an incremental, demand-driven CFA to keep the model synchronized with a changing program.

More recently, Sui and Xue [27] developed SUPA, an on-demand analyzer for C programs which refines value flows during analysis. SUPA is designed for low-budget environments and its evaluation explicitly weighs its ultimate precision against its initial budget. Demand *m*-CFA can be positioned similarly where the low-budget environments are compilers and IDEs for functional programs.

10 Future Work

Having established context sensitivity for Demand CFA, we intend to: (1) develop a general operational framework in which a variety of effects can be expressed, including mutation, (2) investigate the tradeoff between precision and the non-detection of dead code within the analysis (occurring during binding resolution), (3) determine whether *m*-CFA's re-binding policy could be adopted to a demand setting, and whether Demand *m*-CFA could benefit from it, (4) determine how to reuse fixpoint caches between queries while keeping environments as indeterminate as possible to create an incremental analysis [25,18], (5) consider how selective context sensitivity [11] could be realized given the indeterminate environments of our approach.

References

1. Biswas, S.K.: A demand-driven set-based analysis. In: Proceedings of the 24th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages. pp. 372–385. POPL '97, ACM, New York, NY, USA (Jan 1997). <https://doi.org/10.1145/263699.263753>
2. Darais, D., Labich, N., Nguyen, P.C., Van Horn, D.: Abstracting definitional interpreters (functional pearl). Proceedings of the ACM on Programming Languages **1**(ICFP), 12:1–12:25 (Aug 2017). <https://doi.org/10.1145/3110256>
3. Darais, D.C.: Mechanizing Abstract Interpretation. Ph.D. thesis, University of Maryland, College Park, MD, USA (2017)
4. Duesterwald, E., Gupta, R., Soffa, M.L.: A practical framework for demand-driven interprocedural data flow analysis. ACM Transactions on Programming Languages and Systems (TOPLAS) **19**(6), 992–1030 (Nov 1997). <https://doi.org/10.1145/267959.269970>
5. Germane, K., McCarthy, J., Adams, M.D., Might, M.: Demand control-flow analysis. In: Enea, C., Piskac, R. (eds.) Verification, Model Checking, and Abstract Interpretation. Lecture Notes in Computer Science, vol. 11388, pp. 226–246. Springer, Cham (2019). https://doi.org/10.1007/978-3-030-11245-5_11
6. Heintze, N., McAllester, D.: Linear-time subtransitive control flow analysis. In: Proceedings of the ACM SIGPLAN 1997 Conference on Programming Language Design and Implementation. pp. 261–272. PLDI '97, ACM, New York, NY, USA (May 1997). <https://doi.org/10.1145/258915.258939>
7. Heintze, N., Tardieu, O.: Demand-driven pointer analysis. In: Proceedings of the ACM SIGPLAN 2001 Conference on Programming Language Design and Implementation. pp. 24–34. PLDI '01, ACM, New York, NY, USA (Jun 2001). <https://doi.org/10.1145/378795.378802>
8. Horwitz, S., Reps, T., Sagiv, M.: Demand interprocedural dataflow analysis. ACM SIGSOFT Software Engineering Notes **20**(4), 104–115 (1995)
9. Johnson, J.I., Sergey, I., Earl, C., Might, M., Van Horn, D.: Pushdown flow analysis with abstract garbage collection. Journal of Functional Programming **24**(2–3), 218–283 (2014). <https://doi.org/10.1017/S0956796814000100>
10. Leijen, D.: Koka programming language. <https://github.com/koka-lang/koka> (2019), accessed: 2024-05-30
11. Li, Y., Tan, T., Møller, A., Smaragdakis, Y.: A principled approach to selective context sensitivity for pointer analysis. ACM Transactions on Programming Languages and Systems (TOPLAS) **42**(2), 1–40 (2020)
12. Lu, Y., Shang, L., Xie, X., Xue, J.: An incremental points-to analysis with CFL-reachability. In: Jhala, R., De Bosschere, K. (eds.) Compiler Construction. Lecture Notes in Computer Science, vol. 7791, pp. 61–81. Springer, Berlin, Heidelberg (2013). https://doi.org/10.1007/978-3-642-37051-9_4
13. Midtgaard, J., Jensen, T.: A calculational approach to control-flow analysis by abstract interpretation. In: Alpuente, M., Vidal, G. (eds.) Static Analysis. Lecture Notes in Computer Science, vol. 5079, pp. 347–362. Springer, Berlin, Heidelberg (2008). https://doi.org/10.1007/978-3-540-69166-2_23
14. Might, M., Shivers, O.: Improving flow analyses via GCFA: abstract garbage collection and counting. In: Proceedings of the Eleventh ACM SIGPLAN International Conference on Functional Programming. pp. 13–25. ICFP '06, ACM, New York, NY, USA (Sep 2006). <https://doi.org/10.1145/1159803.1159807>

15. Might, M., Smaragdakis, Y., Van Horn, D.: Resolving and exploiting the k -CFA paradox: illuminating functional vs. object-oriented program analysis. In: Proceedings of the 31st ACM SIGPLAN Conference on Programming Language Design and Implementation. pp. 305–315. PLDI '10, ACM, New York, NY, USA (Jun 2010). <https://doi.org/10.1145/1806596.1806631>
16. Nicolay, J., Stiévenart, Q., De Meuter, W., De Roover, C.: Effect-driven flow analysis. In: Verification, Model Checking, and Abstract Interpretation: 20th International Conference, VMCAI 2019, Cascais, Portugal, January 13–15, 2019, Proceedings 20. pp. 247–274. Springer (2019)
17. Palmer, Z., Smith, S.F.: Higher-order demand-driven program analysis. In: Krishnamurthi, S., Lerner, B.S. (eds.) 30th European Conference on Object-Oriented Programming (ECOOP 2016). Leibniz International Proceedings in Informatics (LIPIcs), vol. 56, pp. 19:1–19:25. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik, Dagstuhl, Germany (2016). <https://doi.org/10.4230/LIPIcs.ECOOP.2016.19>
18. Van der Plas, J., Stiévenart, Q., Van Es, N., De Roover, C.: Incremental flow analysis through computational dependency reification. In: 2020 IEEE 20th International Working Conference on Source Code Analysis and Manipulation (SCAM). pp. 25–36 (2020). <https://doi.org/10.1109/SCAM51674.2020.00008>
19. Saha, D., Ramakrishnan, C.R.: Incremental and demand-driven points-to analysis using logic programming. In: Proceedings of the 7th ACM SIGPLAN International Conference on Principles and Practice of Declarative Programming. pp. 117–128. PPDP '05, ACM, New York, NY, USA (Jul 2005). <https://doi.org/10.1145/1069774.1069785>
20. Schoepe, D., Seekatz, D., Stoilkovska, I., Stucki, S., Tattersall, D., Bolignano, P., Raimondi, F., Chang, B.Y.E.: Lifting on-demand analysis to higher-order languages. In: International Static Analysis Symposium. pp. 460–484. Springer (2023)
21. Shang, L., Xie, X., Xue, J.: On-demand dynamic summary-based points-to analysis. In: Proceedings of the Tenth International Symposium on Code Generation and Optimization. pp. 264–274. CGO '12, ACM, New York, NY, USA (Mar 2012). <https://doi.org/10.1145/2259016.2259050>
22. Shivers, O.: Control-Flow Analysis of Higher-Order Languages. Ph.D. thesis, Carnegie Mellon University, Pittsburgh, PA, USA (1991)
23. Späth, J., Do, L.N.Q., Ali, K., Bodden, E.: Boomerang: Demand-driven flow- and context-sensitive pointer analysis for Java. In: Krishnamurthi, S., Lerner, B.S. (eds.) 30th European Conference on Object-Oriented Programming, ECOOP 2016, July 18–22, 2016, Rome, Italy. Leibniz International Proceedings in Informatics (LIPIcs), vol. 56, pp. 22:1–22:26. Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik, Dagstuhl, Germany (2016). <https://doi.org/10.4230/LIPIcs.ECOOP.2016.22>
24. Sperber, M., Dybvig, R.K., Flatt, M., van Straaten, A., Fidler, R., Matthews, J.: Revised [6] Report on the Algorithmic Language Scheme. Cambridge University Press, New York, NY, USA, 1st edn. (Jun 2010)
25. Stein, B., Chang, B.Y.E., Sridharan, M.: Demanded abstract interpretation. In: Proceedings of the 42nd ACM SIGPLAN International Conference on Programming Language Design and Implementation. pp. 282–295 (2021)
26. Su, Y., Ye, D., Xue, J.: Parallel pointer analysis with CFL-reachability. In: 43rd International Conference on Parallel Processing (ICPP 2014). pp. 451–460. IEEE Computer Society, Los Alamitos, CA, USA (2014). <https://doi.org/10.1109/ICPP.2014.54>

27. Sui, Y., Xue, J.: On-demand strong update analysis via value-flow refinement. In: Proceedings of the 2016 24th ACM SIGSOFT International Symposium on Foundations of Software Engineering. pp. 460–473. FSE 2016, ACM, New York, NY, USA (Nov 2016). <https://doi.org/10.1145/2950290.2950296>
28. Whiting, T.: Demand m-cfa koka implementation. <https://github.com/TimWhiting/koka/tree/esop-2025> (2024)
29. Whiting, T.: Demand m-cfa pure scheme implementation. <https://github.com/TimWhiting/context-sensitive-demand-cfa> (2024)
30. Whiting, T., Germane, K.: Demand m-cfa technical report. <https://timwhiting.dev/assets/pdfs/demand-mcfa-techreport.pdf> (2024)

Open Access. This chapter is licensed under the terms of the Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0/>), which permits use, sharing, adaptation, distribution, and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license and indicate if changes were made.

The images or other third party material in this chapter are included in the chapter's Creative Commons license, unless indicated otherwise in a credit line to the material. If material is not included in the chapter's Creative Commons license and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder.

