

Control-Flow Analysis of Function Calls and Returns by Abstract Interpretation

Jan Midtgaard
Roskilde University
jmid@ruc.dk

Thomas P. Jensen
CNRS
thomas.jensen@irisa.fr

Abstract

We derive a control-flow analysis that approximates the interprocedural control-flow of both function calls and returns in the presence of first-class functions and tail-call optimization. In addition to an abstract environment, our analysis computes for each expression an abstract control stack, effectively approximating where function calls return across optimized tail calls. The analysis is systematically calculated by abstract interpretation of the stack-based C_aEK abstract machine of Flanagan et al. using a series of Galois connections. Abstract interpretation provides a unifying setting in which we 1) prove the analysis equivalent to the composition of a continuation-passing style (CPS) transformation followed by an abstract interpretation of a stack-less CPS machine, and 2) extract an equivalent constraint-based formulation, thereby providing a rational reconstruction of a constraint-based control-flow analysis from abstract interpretation principles.

Categories and Subject Descriptors F.3.2 [Logics and Meanings of Programs]: Semantics of Programming Languages—Program Analysis

General Terms Languages, Theory, Verification

Keywords Control flow analysis, abstract interpretation, tail-call optimization, continuation-passing style, direct style, constraint-based analysis

1. Introduction

The control flow of a functional program is expressed in terms of function calls and returns. As a result, iteration in functional programs is expressed using recursive functions. In order for this approach to be feasible, language implementations perform *tail-call optimization* of function calls [Clinger, 1998], by not pushing a stack frame on the control stack at call sites in *tail position*. Consequently functions do not necessarily return control to their caller. Control-flow analysis (CFA) has long been a staple of program optimization and verification. Surprisingly, research on control-flow analysis has focused on calls: A textbook CFA “will determine where the flow of control may be transferred to in the case [...] of a function application.” [Nielsen et al., 1999]. Our systematic approximation of a known operational semantics leads to a CFA

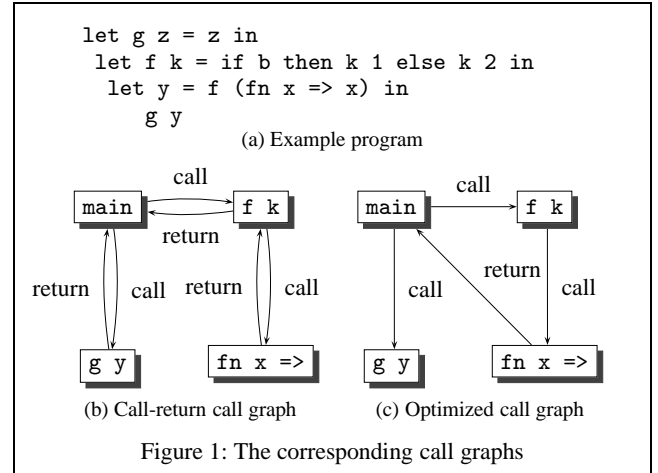


Figure 1: The corresponding call graphs

that “will determine where the flow of control may be transferred to in the case of a function return.” The resulting analysis thereby approximates both call and return information for a higher-order, direct-style language. Interestingly it does so by approximating the control stack.

Consider the example program in Fig. 1(a). The program contains three functions: two named function `g` and `f` and an anonymous function `fn x => x`. A standard direct-style CFA can determine that the applications of `k` in each branch of the conditional will call the anonymous function `fn x => x` at run time. Building a call-graph based on this output gives rise to Fig. 1(b), where we have named the main expression of the program `main`. In addition to the above resolved call, our analysis will determine that the anonymous function returns to the let-binding of `y` in `main` upon completion, rather than to its caller. The analysis hence gives rise to the call graph in Fig. 1(c).

On a methodological level, we derive the analysis systematically by Cousot-Cousot-style *abstract interpretation*. The analysis approximates the reachable states of an existing abstract machine from the literature: the C_aEK machine of Flanagan et al. [1993]. We obtain the analysis as the result of composing the collecting semantics induced by the abstract machine with a series of Galois connections that each specifies one aspect of the abstraction in the analysis.

We show how the abstract interpretation formulation lends itself to a lock-step equivalence proof between our analysis and a previously derived CPS-based CFA. More precisely, we define a relation between the abstract domains of the analyses that is a simulation between the two, reducing the proof to a fixpoint induction over the abstract interpretations.

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

ICFP'09, August 31–September 2, 2009, Edinburgh, Scotland, UK.
Copyright © 2009 ACM 978-1-60558-332-7/09/08...\$10.00

To sum up, the main contributions of this article are:

- An abstract interpretation-derivation of a CFA for a higher-order functional language from a well-known operational semantics,
- a resulting CFA with *reachability* which computes both call and return control-flow,
- a proof of equivalence of the analysis of programs in direct style and the CPS analysis of their CPS counterparts,
- an equivalent constraint-based analysis extracted from the above.

1.1 Related work

We separate the discussion of related analyses in two: direct-style analyses and analyses based on CPS.

Direct-style CFA has a long research history. Jones [1981] initially developed methods for approximating the control flow of lambda terms. Since then Sestoft [1989] conceived the related *closure analysis*. Palsberg [1995] simplified the analysis and formulated an equivalent constraint-based analysis. At the same time Heintze [1994] developed a related set-based analysis formulated in terms of set constraints. For a detailed account of related work, we refer to a recent survey of the area [Midtgaard, 2007]. It is worth emphasizing that all of the above analyses focus on calls, in that they approximate the source lambdas being called at each call site. As such they do not directly determine return flow for programs in direct style.

CPS-based CFA was pioneered by Shivers [1988] who formulated control-flow analysis for Scheme. Since then several analyses have been formulated for CPS [Ayers, 1992, Ashley and Dybvig, 1998, Might and Shivers, 2006]. In CPS all calls are tail calls, and even returns are encoded as calls to the current continuation. By determining “call flow” and hence the receiver functions of such continuation calls, a CPS-based CFA thereby determines return flow without additional effort.

The impact of CPS transformation on static analyses originates in binding-time analysis, for which the transformation is known to have a positive effect [Consel and Danvy, 1991, Damian and Danvy, 2003]. As to the impact of CPS transformation on CFA we separate the previous work on the subject in two:

1. results relating an analysis *specialized* to the source language to an analysis *specialized* to the target language (CPS), and
2. results relating the analysis of a program to the *same analysis* of the CPS transformed program.

Sabry and Felleisen [1994] designed and compared specialized analyses and hence falls into the first category as does the present paper. Damian and Danvy [2003] related the analysis of a program and its CPS counterpart for a standard flow-logic CFA (as well as for two binding-time analyses), and Palsberg and Wand [2003] related the analysis of a program and its CPS counterpart for a standard conditional constraint CFA. Hence the latter two fall into the second category.

We paraphrase the relevant theorems of Sabry and Felleisen [1994], of Damian and Danvy [2003], of Palsberg and Wand [2003], and of the present paper in order to underline the difference between the contributions (C refers to non-trivial, 0-CFA-like analyses defined in the cited papers, p ranges over direct-style programs, cps denotes CPS transformation, and \sim denotes analysis equivalence). Our formulations should not be read as a formal system, but only as a means for elucidating the difference between the contributions.

Sabry and Felleisen [1994]:

exists analyses C_1, C_2 : exists $p, C_1(p) \approx C_2(cps(p))$

Damian and Danvy [2003], Palsberg and Wand [2003]:

exists analysis C : for all $p, C(p) \sim C(cps(p))$

Present paper, Theorem 5.1:

exists analyses C_1, C_2 : for all $p, C_1(p) \sim C_2(cps(p))$

Our work relates to all of the above contributions. The disciplined derivation of specialized CPS and direct-style analyses results in comparable analyses, contrary to Sabry and Felleisen [1994]. Furthermore our equivalence proof extends the results of Damian and Danvy [2003] and Palsberg and Wand [2003] in that we relate both call flow, *return flow*, and *reachability*, contrary to their relating only the call flow of standard CFAs. In addition, the systematic abstract interpretation-based approach suggests a strategy for obtaining similar equivalence results for other CFAs derived in this fashion.

Formulating CFA in the traditional abstract interpretation framework was stated as an open problem by Nielson and Nielson [1997]. It has been a recurring theme in the work of the present authors. In an earlier paper Spoto and Jensen [2003] investigated class analysis of object-oriented programs as a Galois connection-based abstraction of a trace semantics. In a recent article [Midtgaard and Jensen, 2008a], the authors systematically derived a CPS-based CFA from the collecting semantics of a stack-less machine. While investigating how to derive a corresponding direct-style analysis we discovered a mismatch between the computed return information.

As tail calls are identified syntactically, the additional information could also have been obtained by a subsequent analysis after a traditional direct-style CFA. However we view the need for such a subsequent analysis as a strong indication of a mismatch between the direct-style and CPS analysis formulations. Debray and Proebsting [1997] have investigated such a “*return analysis*” for a first-order language with tail-call optimization. This paper builds a semantics-based CFA that determines such information, and for a higher-order language.


The systematic design of constraint-based analyses is a goal shared with the *flow logic* framework of Nielson and Nielson [2002]. In flow logic an analysis specification can be systematically transformed into a constraint-based analysis. The present paper instead extracts a constraint-based analysis from an analysis developed in the original abstract interpretation framework.

The idea of CFA by control stack approximation, applies equally well to imperative or object-oriented programs, but it is beyond the scope of this paper to argue this point. Due to space limitations most calculations and proofs are also omitted. We refer the reader to the accompanying technical report [Midtgaard and Jensen, 2008b].

2. Language and semantics

Our source language is a simple call-by-value core language known as *administrative normal form* (ANF). The grammar of ANF terms is given in Fig. 2(a). Following Reynolds, the grammar distinguishes *serious* expressions, i.e., terms whose evaluation may diverge, from *trivial* expressions, i.e., terms without risk of divergence. Trivial expressions include constants, variables, and functions, and serious expressions include returns, let-bindings, tail calls, and non-tail calls. Programs are serious expressions.

The analysis is calculated from a simple operational semantics in the form of an abstract machine. We use the environment-based C_aEK abstract machine of Flanagan et al. [1993] given in Fig. 2 in which functional values are represented using *closures*, i.e., pairs of a lambda-expression and an environment. The environment-component captures the (values of the) free variables of the lambda. Machine states are triples consisting of a serious expression, an



environment and a control stack. The control stack is composed of elements (“stack frames”) of the form $[x, s, e]$ where x is the variable receiving the return value w of the current function call, and s is a serious expression whose evaluation in the environment $e[x \mapsto w]$ represents the rest of the computation in that stack frame. The empty stack is represented by stop . The machine has a helper function μ for evaluation of trivial expressions. The machine is initialized with the input program, with an empty environment, and with an initial stack, that will bind the result of the program to a special variable x_r before halting. Evaluation follows by repeated application of the machine transitions.

3. Abstract interpretation basics

We assume some familiarity with the basic mathematical facts recalled in Appendix A. Canonical abstract interpretation approximates the *collecting semantics* of a transition system [Cousot, 1981]. A standard example of a collecting semantics is the *reachable states* from a given set of initial states I . Given a transition function T defined as: $T(\Sigma) = I \cup \{\sigma \mid \exists \sigma' \in \Sigma : \sigma' \rightarrow \sigma\}$, we can compute the reachable states of T as the least fixed-point $\text{lfp } T$ of T . The collecting semantics is ideal, in that it is the most precise analysis. Unfortunately it is in general uncomputable. Abstract interpretation therefore approximates the collecting semantics, by instead computing a fixed-point over an alternative and perhaps simpler domain. For this reason, abstract interpretation is also referred to as a theory of fixed-point approximation.

Abstractions are formally represented as Galois connections which connect complete lattices through a pair of adjoint functions α and γ (see Appendix A). Galois connection-based abstract interpretation suggests that one may derive an analysis systematically by composing the transition function with these adjoints: $\alpha \circ T \circ \gamma$. In this setting Galois connections allow us to gradually refine the collecting semantics into a computable analysis function by mere calculation. An alternative “recipe” consists in rewriting the composition of the abstraction function and transition function $\alpha \circ T$ into something of the form $T^\sharp \circ \alpha$, from which the analysis function T^\sharp can be read off [Cousot and Cousot, 1992a]. Cousot [1999] has shown how to systematically construct a static analyser for a first-order imperative language using calculational abstract interpretation.

4. Approximating the C_aEK collecting semantics

As our collecting semantics we consider the reachable states of the C_aEK machine, expressed as the least fixed point $\text{lfp } F$ of the following transition function.

$$F : \wp(C \times Env \times K) \rightarrow \wp(C \times Env \times K)$$

$$F(S) = I_p \cup \{s \mid \exists s' \in S : s' \longrightarrow s\}$$

where $I_p = \{\langle p, \bullet, [x_r, x_r, \bullet] :: \text{stop} \rangle\}$

First we formulate in Fig. 3(a) an equivalent helper function μ_c extended to work on sets of environments.

Lemma 4.1. $\forall t, e : \{\mu(t, e)\} = \mu_c(t, \{e\})$

The equivalence of the two helper functions follow straightforwardly. This lemma enables us to express an equivalent collecting semantics based on μ_c , which appears in Fig. 3. The equivalence of F and F^c follows from the lemma and by unfolding the definitions.

The abstraction of the collecting semantics is staged in several steps. Figure 4 provides an overview. Intuitively, the analysis extracts three pieces of information from the set of reachable states.

1. An approximation of the set of reachable expressions.
2. A relation between expressions and control stacks that represents where the values of expressions are returned to.
3. An abstract environment mapping variables to the expressions that may be bound to that variable. This is standard in CFA and allows to determine which functions are called at a given call site.

Keeping an explicit set of reachable expressions is more precise than leaving it out, once we further approximate the expression-stack pairs. Alternatively the reachable expressions would be approximated by the expressions present in the expression-stack relation. However expressions may be in the expression-stack relation without ever being reached. An example hereof would be a diverging non-tail call.

To formalize this intuition, we first perform a Cartesian abstraction of the machine states, however keeping the relation between expressions and their corresponding control stacks. The second step in the approximation consists in closing the triples by a closure operator, to ensure that (a) any saved environment on the stack or nested within another environment is itself part of the environment

$$\begin{aligned} \mu_c : T \times \wp(\text{Env}) &\rightarrow \wp(\text{Val}) \\ \mu_c(c, E) &= \{c\} \\ \mu_c(x, E) &= \{w \mid \exists e \in E : w = e(x)\} \\ \mu_c(\text{fn } x \Rightarrow s, E) &= \{\{\text{fn } x \Rightarrow s, e\} \mid \exists e \in E\} \end{aligned}$$

(a) Helper function

$$\begin{aligned} F^c : \wp(C \times \text{Env} \times K) &\rightarrow \wp(C \times \text{Env} \times K) \\ F^c(S) = I_p & \\ \cup \bigcup_{\substack{\langle t, e, [x, s', e'] :: k \rangle \in S \\ w \in \mu_c(t, \{e\})}} \{ \langle s', e'[x \mapsto w], k' \rangle \} & \\ \cup \bigcup_{\substack{\langle \text{let } x = t \text{ in } s, e, k \rangle \in S \\ w \in \mu_c(t, \{e\})}} \{ \langle s, e[x \mapsto w], k \rangle \} & \\ \cup \bigcup_{\substack{\langle t_0 t_1, e, k \rangle \in S \\ [\text{fn } x \Rightarrow s', e'] \in \mu_c(t_0, \{e\}) \\ w \in \mu_c(t_1, \{e\})}} \{ \langle s', e'[x \mapsto w], k' \rangle \} & \\ \cup \bigcup_{\substack{\langle \text{let } x = t_0 t_1, \text{ in } s, e, k \rangle \in S \\ [\text{fn } y \Rightarrow s', e'] \in \mu_c(t_0, \{e\}) \\ w \in \mu_c(t_1, \{e\})}} \{ \langle s', e'[y \mapsto w], [x, s, e] :: k \rangle \} & \end{aligned}$$

(b) Transition function

Figure 3: Collecting semantics

$\wp(C \times \text{Env} \times K)$	coll. sem.	F^c
$\alpha_x \updownarrow \gamma_x$		
$\wp(C) \times \wp(C \times K) \times \wp(\text{Env})$	-	F^\times
$\rho \updownarrow 1$		
$\rho(\wp(C) \times \wp(C \times K) \times \wp(\text{Env}))$	-	F^ρ
$\alpha_\otimes \updownarrow \gamma_\otimes$		
$\wp(C) \times (C/\equiv \rightarrow \wp(K^\sharp)) \times \text{Env}^\sharp$	0-CFA	F^\sharp

Figure 4: Overview of abstraction

set, and (b) that all expression-control stack pairs that appear further down in a control stack are also contained in the expression-stack relation. We explain this in more detail below (Section 4.2). Finally as a third step we approximate stacks by their top element, we merge expressions with the same return point into equivalence classes, and we approximate closure values by their lambda expression.

In the following sections we provide a detailed explanation of each abstraction in turn.

4.1 Projecting machine states

The mapping that extracts the three kinds of information described above is defined formally as follows.

$$\begin{aligned} \wp(C \times \text{Env} \times K) &\xrightarrow[\alpha_x]{\gamma_x} \wp(C) \times \wp(C \times K) \times \wp(\text{Env}) \\ \alpha_x(S) &= \langle \pi_1 S, \{ \langle s, k \rangle \mid \exists e : \langle s, e, k \rangle \in S \}, \pi_2 S \rangle \\ \gamma_x(\langle C, F, E \rangle) &= \{ \langle s, e, k \rangle \mid s \in C \wedge \langle s, k \rangle \in F \wedge e \in E \} \end{aligned}$$

Lemma 4.2. α_x, γ_x is a Galois connection.

The above Galois connection and the proof hereof closely resembles the independent attributes abstraction, which is a known Galois connection. We use the notation \cup_x and \subseteq_x for the componentwise join and componentwise inclusion of triples.

As traditional [Cousot and Cousot, 1979, 1992a, 1994], we will assume that the abstract product domains throughout this article have been *reduced*, i.e., all triples $\langle A, B, C \rangle$ with a bottom component ($A = \perp_a \vee B = \perp_b \vee C = \perp_c$) have been eliminated and replaced by a single bottom element $\langle \perp_a, \perp_b, \perp_c \rangle$.

Based on this abstraction we can now calculate a new transfer function F^\times . The resulting transition function appears in Fig. 5. By construction, the transition function satisfies the following theorem.

Theorem 4.1.

$$\forall C, F, E : \alpha_x(F^c(\gamma_x(\langle C, F, E \rangle))) = F^\times(\langle C, F, E \rangle)$$

4.2 A closure operator on machine states

For the final analysis, we are only interested in an abstraction of the information present in an expression-stack pair. More precisely, we aim at only keeping track of the link between an expression and the top stack frame in effect during its evaluation, throwing away everything below. However, we need to make this information explicit for all expressions appearing on the control stack, i.e., for a pair $\langle s, [x, s', e] :: k \rangle$ we also want to retain that s' will be evaluated with control stack k . Similarly, environments can be stored on the stack or inside other environments and will have to be extracted. We achieve this by defining a suitable *closure operator* on these nested structures.

For environments, we adapt the definition of a constituent relation due to Milner and Tofte [1991]. We say that each component x_i of a tuple $\langle x_0, \dots, x_n \rangle$ is a *constituent* of the tuple, written $\langle x_0, \dots, x_n \rangle \succ x_i$. For a partial function¹ $f = [x_0 \mapsto w_0, \dots, x_n \mapsto w_n]$, we say that each w_i is a constituent of the function, written $f \succ w_i$. We write \succ^* for the reflexive, transitive closure of the constituent relation.

To deal with the control stack, we define an order on expression-stack pairs. Two pairs are ordered if (a) the stack component of the second is the tail of the first's stack component, and (b) the expression component of the second, resides on the top stack frame of the first pair: $\langle s, [x, s', e] :: k \rangle \succ \langle s', k \rangle$. We write \succ^* for the reflexive, transitive closure of the expression-stack pair ordering.

Next, we consider an operator ρ , defined in terms of the constituent relation and the expression-stack pair ordering. The operator ρ ensures that all constituent environments will themselves belong to the set of environments, and that any structurally smaller expression-stack pairs are also contained in the expression-stack relation.

Definition 4.1.

$$\begin{aligned} \rho(\langle C, F, E \rangle) &= \langle C, \{ \langle s, k \rangle \mid \exists \langle s', k' \rangle \in F : \langle s', k' \rangle \succ^* \langle s, k \rangle, \\ &\quad \{ e \mid \exists \langle s, k \rangle \in F : \langle s, k \rangle \succ^* e \vee \exists e' \in E : e' \succ^* e \} \rangle \end{aligned}$$

We need to relate the expression-stack ordering to the constituent relation. By case analysis one can prove that $\forall \langle s, k \rangle, \langle s', k' \rangle : \langle s, k \rangle \succ \langle s', k' \rangle \implies k \succ k'$. By structural induction (on the stack component) it now follows that $\forall \langle s, k \rangle, \langle s', k' \rangle : \langle s, k \rangle \succ^* \langle s', k' \rangle \implies k \succ^* k'$. Based on these results we can verify that ρ is a closure operator and formulate an abstraction on the triples:

$$\wp(C) \times \wp(C \times K) \times \wp(\text{Env}) \xrightarrow[\rho]{1} \rho(\wp(C) \times \wp(C \times K) \times \wp(\text{Env}))$$

¹ Milner and Tofte define the constituent relation for finite functions.

$$\begin{aligned}
F^\times &: \wp(C) \times \wp(C \times K) \times \wp(Env) \rightarrow \wp(C) \times \wp(C \times K) \times \wp(Env) \\
F^\times(\langle C, F, E \rangle) &= \langle \{p\}, \{ \langle p, [x_r, x_r, \bullet] :: \text{stop} \rangle \}, \{\bullet\} \rangle \\
&\cup_x \bigcup_x \langle \{s'\}, \{ \langle s', k' \rangle \}, \{e' [x \mapsto w]\} \rangle \\
&\langle \{t\}, \{ \langle t, [x, s', e'] :: k' \rangle \}, \{e\} \rangle_{\subseteq_x \langle C, F, E \rangle} \\
&\quad w \in \mu_c(t, \{e\}) \\
&\cup_x \bigcup_x \langle \{s\}, \{ \langle s, k \rangle \}, \{e [x \mapsto w]\} \rangle \\
&\langle \{ \text{let } x=t \text{ in } s \}, \{ \langle \text{let } x=t \text{ in } s, k \rangle \}, \{e\} \rangle_{\subseteq_x \langle C, F, E \rangle} \\
&\quad w \in \mu_c(t, \{e\}) \\
&\cup_x \bigcup_x \langle \{s'\}, \{ \langle s', k \rangle \}, \{e' [x \mapsto w]\} \rangle \\
&\langle \{t_0 t_1\}, \{ \langle t_0 t_1, k \rangle \}, \{e\} \rangle_{\subseteq_x \langle C, F, E \rangle} \\
&\quad [\text{fn } x \Rightarrow s', e'] \in \mu_c(t_0, \{e\}) \\
&\quad w \in \mu_c(t_1, \{e\}) \\
&\cup_x \bigcup_x \langle \{s'\}, \{ \langle s', [x, s, e] :: k \rangle \}, \{e' [y \mapsto w]\} \rangle \\
&\langle \{ \text{let } x=t_0 t_1 \text{ in } s \}, \{ \langle \text{let } x=t_0 t_1 \text{ in } s, k \rangle \}, \{e\} \rangle_{\subseteq_x \langle C, F, E \rangle} \\
&\quad [\text{fn } y \Rightarrow s', e'] \in \mu_c(t_0, \{e\}) \\
&\quad w \in \mu_c(t_1, \{e\})
\end{aligned}$$

Figure 5: Abstract transition function

We use the notation \cup_ρ for the join operation $\lambda X. \rho(\cup_x X)$ on the closure operator-induced complete lattice. First observe that in our case:

$$\cup_\rho = \lambda X. \rho(\bigcup_x X_i) = \lambda X. \bigcup_x \rho(X_i) = \lambda X. \bigcup_x X_i = \cup_x$$

Based on the closure operator-based Galois connection, we can calculate a new intermediate transfer function F^ρ . The resulting transfer function appears in Fig. 6. This transfer function differs only minimally from the one in Fig. 5, in that (a) the signature has changed, (b) the set of initial states has been “closed” and now contains the structurally smaller pair $\langle x_r, \text{stop} \rangle$, and (c) the four indexed joins now each join “closed” triples in the image of the closure operator. By construction, the new transition function satisfies the following theorem.

Theorem 4.2. $\forall C, F, E: \rho \circ F^\times \circ 1(\langle C, F, E \rangle) = F^\rho(\langle C, F, E \rangle)$

4.3 Abstracting the expression-stack relation

Since stacks can grow unbounded (for non-tail recursive programs), we need to approximate the stack component and hereby the expression-stack relation. We first formulate a grammar of abstract stacks and an elementwise operator $@ : C \times K \rightarrow C \times K^\sharp$ operating on expression-stack pairs.

$$\begin{aligned}
K^\sharp \ni k^\sharp &::= \text{stop} \mid [x, s] \\
@(\langle s, \text{stop} \rangle) &= \langle s, \text{stop} \rangle \\
@(\langle s, [x, s', e] :: k \rangle) &= \langle s, [x, s'] \rangle
\end{aligned}$$

Based on the elementwise operator we can now use an elementwise abstraction.

Elementwise abstraction [Cousot and Cousot, 1997]: A given elementwise operator $@ : C \rightarrow A$ induces a Galois connection:

$$\begin{aligned}
\langle \wp(C); \subseteq \rangle &\xleftrightarrow[\alpha_{@}]{\gamma_{@}} \langle \wp(A); \subseteq \rangle \\
\alpha_{@}(P) &= \{ @ (p) \mid p \in P \} \quad \gamma_{@}(Q) = \{ p \mid @ (p) \in Q \}
\end{aligned}$$

Notice how some expressions share the same return point (read: same stack): the expressions $\text{let } x=t \text{ in } s$ and s share the same return point, and $\text{let } x=t_0 t_1 \text{ in } s$ and s share the same return point. In order to eliminate such redundancy we define an equivalence relation on serious expressions grouping together expressions sharing

the same return point. We define the smallest equivalence relation \equiv satisfying:

$$\begin{aligned}
\text{let } x=t \text{ in } s &\equiv s \\
\text{let } x=t_0 t_1 \text{ in } s &\equiv s
\end{aligned}$$

Based hereon we define a second elementwise operator $@' : C \times K^\sharp \rightarrow C/\equiv \times K^\sharp$ mapping the first component of an expression-stack pair to a representative of its corresponding equivalence class:

$$@'(\langle s, k^\sharp \rangle) = \langle [s]_{\equiv}, k^\sharp \rangle$$

We can choose the outermost expression as a representative for each equivalence class by a linear top-down traversal of the input program.

Pointwise coding of a relation [Cousot and Cousot, 1994]: A relation can be isomorphically encoded as a set-valued function by a Galois connection:

$$\langle \wp(A \times B); \subseteq \rangle \xleftrightarrow[\alpha_\omega]{\gamma_\omega} \langle A \rightarrow \wp(B); \subseteq \rangle$$

$$\alpha_\omega(r) = \lambda a. \{ b \mid \langle a, b \rangle \in r \} \quad \gamma_\omega(f) = \{ \langle a, b \rangle \mid b \in f(a) \}$$

By composing the three above Galois connections we obtain our abstraction of the expression-stack relation:

$$\wp(C \times K) \xleftrightarrow[\alpha_{st}]{\gamma_{st}} C/\equiv \rightarrow \wp(K^\sharp)$$

where $\alpha_{st} = \alpha_\omega \circ \alpha_{@'} \circ \alpha_{@} = \lambda F. \bigcup_{\langle s, k \rangle \in F} \alpha_\omega(\{ @' \circ @(\langle s, k \rangle) \})$ and $\gamma_{st} = \gamma_{@} \circ \gamma_{@'} \circ \gamma_\omega$. We can now prove a lemma relating the concrete and abstract expression-stack relations.

Lemma 4.3. Control stack and saved environments

Let $\langle C, F, E \rangle \in \rho(\wp(C) \times \wp(C \times K) \times \wp(Env))$ be given.

$$\begin{aligned}
\langle s, [x, s', e] :: k \rangle \in F &\implies e \in E \wedge \{ \langle s', k \rangle \} \subseteq F \\
&\wedge \{ [x, s'] \} \subseteq \alpha_{st}(F)([s]_{\equiv})
\end{aligned}$$

Proof. The first half follows from the assumptions. The second half follows from monotonicity of α_{st} , and the definitions of α_{st} , \bigcup , $@$, $@'$, α_ω , and \subseteq . \square

4.4 Abstracting environments

We also abstract values using an elementwise abstraction. Again we formulate a grammar of abstract values and an elementwise

$$\begin{aligned}
& F^P : \rho(\wp(C) \times \wp(C \times K) \times \wp(Env)) \rightarrow \rho(\wp(C) \times \wp(C \times K) \times \wp(Env)) \\
F^P(\langle C, F, E \rangle) = & \langle \{p\}, \{p, [x_r, x_r, \bullet] :: \text{stop}\}, \langle x_r, \text{stop} \rangle, \{\bullet\} \rangle \\
& \cup_x \bigcup_x \rho(\langle \{s'\}, \{s', k'\}, \{e'[x \mapsto w]\} \rangle) \\
& \langle \{t\}, \{t, [x, s', e'] :: k'\}, \{e\} \rangle \subseteq_x \langle C, F, E \rangle \\
& \quad w \in \mu_c(t, \{e\}) \\
& \cup_x \bigcup_x \rho(\langle \{s\}, \{s, k\}, \{e[x \mapsto w]\} \rangle) \\
& \langle \{\text{let } x=t \text{ in } s\}, \{\text{let } x=t \text{ in } s, k\}, \{e\} \rangle \subseteq_x \langle C, F, E \rangle \\
& \quad w \in \mu_c(t, \{e\}) \\
& \cup_x \bigcup_x \rho(\langle \{s'\}, \{s', k\}, \{e'[x \mapsto w]\} \rangle) \\
& \langle \{t_0 t_1\}, \{t_0 t_1, k\}, \{e\} \rangle \subseteq_x \langle C, F, E \rangle \\
& \quad [\text{fn } x \Rightarrow s', e'] \in \mu_c(t_0, \{e\}) \\
& \quad w \in \mu_c(t_1, \{e\}) \\
& \cup_x \bigcup_x \rho(\langle \{s'\}, \{s', [x, s, e] :: k\}, \{e'[y \mapsto w]\} \rangle) \\
& \langle \{\text{let } x=t_0 t_1 \text{ in } s\}, \{\text{let } x=t_0 t_1 \text{ in } s, k\}, \{e\} \rangle \subseteq_x \langle C, F, E \rangle \\
& \quad [\text{fn } y \Rightarrow s', e'] \in \mu_c(t_0, \{e\}) \\
& \quad w \in \mu_c(t_1, \{e\})
\end{aligned}$$

Figure 6: The second abstract transition function

operator $@ : Val \rightarrow Val^\sharp$ mapping concrete to abstract values.

$$\begin{aligned}
Val^\sharp \ni w^\sharp & ::= c \mid [\text{fn } x \Rightarrow s] \\
@(c) & = c \\
@([\text{fn } x \Rightarrow s, e]) & = [\text{fn } x \Rightarrow s]
\end{aligned}$$

The abstraction of environments, which are partial functions, can be composed by a series of well-known Galois connections.

Pointwise abstraction of a set of functions [Cousot and Cousot, 1994]: A given Galois connection on the co-domain $\langle \wp(C); \sqsubseteq \rangle$ $\xrightarrow[\alpha]{\gamma} \langle C^\sharp; \sqsubseteq \rangle$ induces a Galois connection on a set of functions:

$$\begin{aligned}
\langle \wp(D \rightarrow C); \sqsubseteq \rangle & \xleftarrow[\alpha_\Pi]{\gamma_\Pi} \langle D \rightarrow C^\sharp; \sqsubseteq \rangle \\
\alpha_\Pi(F) & = \lambda d. \alpha(\{f(d) \mid f \in F\}) \\
\gamma_\Pi(A) & = \{f \mid \forall d : f(d) \in \gamma(A(d))\}
\end{aligned}$$

Subset abstraction [Cousot and Cousot, 1997]: Given a set C and a strict subset $A \subset C$ hereof, the restriction to the subset induces a Galois connection:

$$\begin{aligned}
\langle \wp(C); \sqsubseteq \rangle & \xleftarrow[\alpha_c]{\gamma_c} \langle \wp(A); \sqsubseteq \rangle \\
\alpha_c(X) & = X \cap A \quad \gamma_c(Y) = Y \cup (C \setminus A)
\end{aligned}$$

A standard trick is to think of partial functions $r : D \rightarrow C$ as total functions $r_\perp : D \rightarrow (C \cup \perp)$ where $\perp \sqsubseteq c$, for all $c \in C$. Consider environments $e \in Var \rightarrow Val$ to be total functions $Var \rightarrow (Val \cup \perp)$ using this idea. In this context the bottom element \perp will denote variable lookup failure. Now compose a subset abstraction $\wp(Val \cup \perp) \xleftarrow[\alpha_c]{\gamma_c} \wp(Val)$ with the above value abstraction, and feed the result to the pointwise abstraction above. The result is a pointwise abstraction of a set of environments, not explicitly modelling variable lookup failure: $\wp(Env) \xleftarrow[\alpha_\Pi]{\gamma_\Pi} Var \rightarrow \wp(Val^\sharp)$. By considering only closed programs, we statically ensure against failure of variable-lookup, hence disregarding \perp loses no information.

4.5 Abstracting the helper function

We can calculate an abstract helper function, by “pushing α ’s” under the function definition, and reading off a resulting abstract definition.

Lemma 4.4. $\forall t, E : \alpha_\circ(\mu_c(t, E)) = \mu^\sharp(t, \alpha_\Pi(E))$

The resulting helper function $\mu^\sharp : T \times Env^\sharp \rightarrow \wp(Val^\sharp)$ reads:

$$\begin{aligned}
\mu^\sharp(c, E^\sharp) & = \{c\} \\
\mu^\sharp(x, E^\sharp) & = E^\sharp(x) \\
\mu^\sharp(\text{fn } x \Rightarrow s, E^\sharp) & = \{[\text{fn } x \Rightarrow s]\}
\end{aligned}$$

where we write Env^\sharp as shorthand for $Var \rightarrow \wp(Val^\sharp)$. We shall need a lemma relating the two helper function definitions on closed environments.

Lemma 4.5. Helper function on closed environments (1) Let $\langle C, F, E \rangle \in \rho(\wp(C) \times \wp(C \times K) \times \wp(Env))$ be given.

$$\begin{aligned}
\{[\text{fn } x \Rightarrow s, e]\} \subseteq \mu_c(t, E) & \implies e \in E \\
\wedge \{[\text{fn } x \Rightarrow s]\} \subseteq \mu^\sharp(t, \alpha_\Pi(E))
\end{aligned}$$

The above lemma is easily extended to capture nested environments in all values returned by the helper function:

Lemma 4.6. Helper function on closed environments (2) Let $\langle C, F, E \rangle \in \rho(\wp(C) \times \wp(C \times K) \times \wp(Env))$ be given.

$$\{w\} \subseteq \mu_c(t, E) \wedge w \succ^* e'' \implies e'' \in E$$

4.6 Abstracting the machine states

We abstract the triplet of sets into abstract triples by a componentwise abstraction.

Componentwise abstraction [Cousot and Cousot, 1994]: Assuming a series of Galois connections: $\wp(C_i) \xleftarrow[\alpha_i]{\gamma_i} A_i$ for $i \in \{1, \dots, n\}$, their componentwise composition induces a Galois connection on tuples:

$$\begin{aligned}
\langle \wp(C_1) \times \dots \times \wp(C_n); \subseteq_x \rangle & \xleftarrow[\alpha_\otimes]{\gamma_\otimes} \langle A_1 \times \dots \times A_n; \subseteq_\otimes \rangle \\
\alpha_\otimes(\langle X_1, \dots, X_n \rangle) & = \langle \alpha_1(X_1), \dots, \alpha_n(X_n) \rangle \\
\gamma_\otimes(\langle x_1, \dots, x_n \rangle) & = \langle \gamma_1(x_1), \dots, \gamma_n(x_n) \rangle
\end{aligned}$$

We write \cup_\otimes and \subseteq_\otimes for componentwise join and inclusion, respectively.

For the set of expressions $\wp(C)$ we use the identity abstraction consisting of two identity functions. For the expression-stack

$$\begin{aligned}
F^\sharp : P &\rightarrow \wp(C) \times (C/\equiv \rightarrow \wp(K^\sharp)) \times Env^\sharp \rightarrow \wp(C) \times (C/\equiv \rightarrow \wp(K^\sharp)) \times Env^\sharp \\
F_p^\sharp(\langle C, F^\sharp, E^\sharp \rangle) &= \langle \{p\}, [[p]_\equiv \mapsto \{[x_r, x_r]\}, [x_r]_\equiv \mapsto \{\text{stop}\}], \lambda_. 0 \rangle \\
&\cup_{\otimes} \bigcup_{\substack{\{t\} \subseteq C \\ \{[x, s']\} \subseteq F^\sharp(t)_\equiv}} \langle \{s'\}, F^\sharp, E^\sharp \cup [x \mapsto \mu^\sharp(t, E^\sharp)] \rangle \\
&\cup_{\otimes} \bigcup_{\{\text{let } x=t \text{ in } s\} \subseteq C} \langle \{s\}, F^\sharp, E^\sharp \cup [x \mapsto \mu^\sharp(t, E^\sharp)] \rangle \\
&\cup_{\otimes} \bigcup_{\substack{\{t_0 t_1\} \subseteq C \\ \{[fn x \Rightarrow s']\} \in \mu^\sharp(t_0, E^\sharp)}} \langle \{s'\}, F^\sharp \cup [[s']_\equiv \mapsto F^\sharp([t_0 t_1]_\equiv)], E^\sharp \cup [x \mapsto \mu^\sharp(t_1, E^\sharp)] \rangle \\
&\cup_{\otimes} \bigcup_{\substack{\{\text{let } x=t_0 t_1 \text{ in } s\} \subseteq C \\ \{[fn y \Rightarrow s']\} \in \mu^\sharp(t_0, E^\sharp)}} \langle \{s'\}, F^\sharp \cup [[s']_\equiv \mapsto \{[x, s]\}], E^\sharp \cup [y \mapsto \mu^\sharp(t_1, E^\sharp)] \rangle
\end{aligned}$$

Figure 7: The resulting analysis function

relation $\wp(C \times K)$ we use the expression-stack abstraction α_{st} developed in Section 4.3. For the set of environments $\wp(Env)$ we use the environment abstraction α_Π developed in Section 4.4.

Using the alternative “recipe” we can calculate the analysis by “pushing α ’s” under the intermediate transition function: $\alpha_{\otimes}(F^P(\langle C, F, E \rangle)) \subseteq_{\otimes} F^\sharp(\langle C, \alpha_{st}(F), \alpha_\Pi(E) \rangle)$ from which the final definition of F^\sharp can be read off. The resulting analysis appears in Fig. 7. The alert reader may have noticed that this final abstraction is not *complete* in that the above equation contains an inequality. Completeness is a desirable goal in an abstract interpretation but unfortunately it is not possible in general without refining the abstract domain [Giacobazzi et al., 2000]. Consider for example the addition operator over the standard *sign-domain*: $0 = \alpha(1 + (-1)) \subseteq \alpha(1) + \alpha(-1) = \top$. As traditional [Cousot, 1999], we instead limit upward judgements to a minimum.

As a corollary of the construction, the analysis safely approximates the reachable states of the abstract machine.

Corollary 4.1. $\alpha_{\otimes} \circ \rho \circ \alpha_{\times}(\text{lfp } F) \subseteq_{\otimes} \text{lfp } F^\sharp$

4.7 Characteristics of the analysis

First of all the analysis incorporates *reachability*: it computes an approximate set of reachable expressions and will only analyse those reachable program fragments. Reachability analyses have previously been discovered independently [Ayers, 1992, Palsberg and Schwartzbach, 1995, Gasser et al., 1997]. In our case they arise naturally from a projecting abstraction of a reachable states collecting semantics.

Second the formulation materializes *monomorphism* into two mappings: (a) one mapping merging all bindings to the same variable, and (b) one mapping merging all calling contexts of the same function. Both characteristics are well known, but our presentation is novel in that it literally captures this phenomenon in two approximation functions.

Third the analysis handles returns inside-out (“*callee-restore*”), in that the called function restores control from the approximate control stack and propagates the obtained return values. This differs from the traditional presentations [Palsberg, 1995, Nielson et al., 1999] that handle returns outside-in (“*caller-restore*”) where the caller propagates the obtained return values from the body of the function to the call site (typically formulated as *conditional constraints*).

$CProg \ni p ::= fn k \Rightarrow e$	(CPS programs)
$SExp \ni e ::= t_0 t_1 c \mid ct$	(serious CPS expressions)
$TExp \ni t ::= x \mid v \mid fn x, k \Rightarrow e$	(trivial CPS expressions)
$CExp \ni c ::= fn v \Rightarrow e \mid k$	(continuation expressions)

Figure 8: BNF of CPS language

5. Analysis equivalence

In previous work [Midtgaard and Jensen, 2008a] we derived an initial CFA with reachability for a CPS language from the stack-less *CE*-machine [Flanagan et al., 1993]. In this section we show that the present ANF analysis achieves the same precision as obtained by first transforming a program into CPS and then using the CPS analysis. This is done by defining a relation that captures how the direct-style analysis and the CPS analysis operate in lock-step.

The grammar of CPS terms is given in Fig. 8. The grammar distinguishes variables in the original source program $x \in X$, from intermediate variables $v \in V$ and continuation variables $k \in K$. We assume the three classes are non-overlapping. Their union constitute the domain of CPS variables $Var = X \cup V \cup K$.

5.1 CPS transformation and back again

In order to state the relation between the ANF and CPS analyses we first recall the relevant program transformations. The below presentation is based on Danvy [1991], Flanagan et al. [1993], and Sabry and Felleisen [1994].

The CPS transformation given in Fig. 9(a) is defined by two mutually recursive functions — one for serious and trivial expressions. A continuation variable k is provided in the initial call to \mathcal{F} . A fresh k is generated in \mathcal{V} ’s lambda abstraction case. To ease the expression of the relation, we choose k unique to the serious expression s — k_s . It follows that we only need one k per lambda abstraction in the original program + an additional k in the initial case.

It is immediate from the definition of \mathcal{F} that the CPS transformation of a let-binding $\text{let } x=t \text{ in } s$ and the CPS transformation of its body s share the same continuation identifier — and similarly for non-tail calls. Hence we shall equate the two:

Definition 5.1. $k_s \equiv k_{s'}$ iff $s \equiv s'$

$\begin{aligned} \mathcal{C} &: P \rightarrow CProg \\ \mathcal{C}[p] &= \text{fn } k_p \Rightarrow \mathcal{F}_{k_p}[p] \\ \mathcal{F} &: K \rightarrow C \rightarrow SExp \\ \mathcal{F}_k[t] &= k \mathcal{V}[t] \\ \mathcal{F}_k[\text{let } x=t \text{ in } s] &= (\text{fn } x \Rightarrow \mathcal{F}_k[s]) \mathcal{V}[t] \\ \mathcal{F}_k[t_0 t_1] &= \mathcal{V}[t_0] \mathcal{V}[t_1] k \\ \mathcal{F}_k[\text{let } x=t_0 t_1 \text{ in } s] &= \mathcal{V}[t_0] \mathcal{V}[t_1] (\text{fn } x \Rightarrow \mathcal{F}_k[s]) \\ \mathcal{V} &: T \rightarrow TExp \\ \mathcal{V}[x] &= x \\ \mathcal{V}[\text{fn } x \Rightarrow s] &= \text{fn } x, k_s \Rightarrow \mathcal{F}_{k_s}[s] \end{aligned}$ <p style="text-align: center;">(a) CPS transformation</p>	$\begin{aligned} \mathcal{D} &: CProg \rightarrow P \\ \mathcal{D}[\text{fn } k \Rightarrow e] &= \mathcal{U}[e] \\ \mathcal{U} &: SExp \rightarrow C \\ \mathcal{U}[k t] &= \mathcal{P}[t] \\ \mathcal{U}[(\text{fn } v \Rightarrow e) t] &= \text{let } v = \mathcal{P}[t] \text{ in } \mathcal{U}[e] \\ \mathcal{U}[t_0 t_1 k] &= \mathcal{P}[t_0] \mathcal{P}[t_1] \\ \mathcal{U}[t_0 t_1 (\text{fn } v \Rightarrow e)] &= \text{let } v = \mathcal{P}[t_0] \mathcal{P}[t_1] \text{ in } \mathcal{U}[e] \\ \mathcal{P} &: TExp \rightarrow T \\ \mathcal{P}[x] &= x \\ \mathcal{P}[v] &= v \\ \mathcal{P}[\text{fn } x, k \Rightarrow e] &= \text{fn } x \Rightarrow \mathcal{U}[e] \end{aligned}$ <p style="text-align: center;">(b) Direct-style transformation</p>
---	---

Figure 9: Transformations to and from CPS

The direct-style transformation given in Fig. 9(b) is defined by two mutually recursive functions over serious and trivial CPS expressions. We define the direct-style transformation of a program $\text{fn } k \Rightarrow e$ as the direct-style transformation of its body $\mathcal{U}[e]$.

Transforming a program, a serious expression, or a trivial expression to CPS and back to direct style yields the original expression, which can be confirmed by (mutual) structural induction on trivial and serious expressions.

Lemma 5.1. $\mathcal{D}[\mathcal{C}[p]] = p \wedge \mathcal{U}[\mathcal{F}_k[s]] = s \wedge \mathcal{P}[\mathcal{V}[t]] = t$

5.2 CPS analysis

We recall the CPS analysis of Midtgaard and Jensen [2008a] in Fig. 10. It is defined as the least fixed point of a program specific transfer function T_p^\sharp . The definition relies on two helper functions μ_r^\sharp and μ_c^\sharp for trivial and continuation expressions, respectively. The analysis computes a pair consisting of (a) a set of serious expressions (the reachable expressions) and (b) an abstract environment. Abstract environments map variables to abstract values. Abstract values can be either the initial continuation stop , function closures $[\text{fn } x, k \Rightarrow e]$, or continuation closures $[\text{fn } v \Rightarrow e]$.

The definition relies on two special variables k_r and v_r , the first of which names the initial continuation and the second of which names the result of the program. To ensure the most precise analysis result, variables in the source program can be renamed to be distinct as is traditional in control-flow analysis [Nielson et al., 1999].

5.3 Analysis equivalence

Before formally stating the equivalence of the two analyses we will study an example run. As our example we use the ANF program: $\text{let } f = \text{fn } x \Rightarrow x \text{ in let } a_1 = f \text{ cn1 in let } a_2 = f \text{ cn2 in } a_2$, taken from Sabry and Felleisen [1994] where we have Church encoded the integer literals. We write cn1 for $\text{fn } s \Rightarrow \text{fn } z \Rightarrow s z$ and cn2 for $\text{fn } s \Rightarrow \text{fn } z \Rightarrow \text{let } t_1 = s z \text{ in } s t_1$. The analysis trace appears in the left column of Table 1.

Similarly we study the CPS analysis of the CPS transformed program. The analysis trace appears in the right column of Table 1 where we have written ccn1 for $\mathcal{V}[\text{cn1}]$ and ccn2 for $\mathcal{V}[\text{cn2}]$. Contrary to Sabry and Felleisen [1994] both the ANF and the CPS analyses achieve the same precision on the example, determining that a_1 will be bound to one of the two integer literals.

We are now in position to state our main theorem relating the ANF analysis to the CPS analysis. Intuitively the theorem relates:

- reachability in ANF to CPS reachability
- abstract stacks in ANF to CPS continuation closures
- abstract stack bottom in ANF to CPS initial continuation
- ANF closures to CPS function closures

Theorem 5.1. Let p be given. Let $\langle C, F^\sharp, E^\sharp \rangle = \text{lfp } F_p^\sharp$ and $\langle Q^\sharp, R^\sharp \rangle = \text{lfp } T_{\mathcal{C}[p]}^\sharp$. Then

$$\begin{aligned} s \in C &\iff \mathcal{F}_{k_s}[s] \in Q^\sharp \wedge \\ [x, s'] \in F^\sharp(s) &\iff [\text{fn } x \Rightarrow \mathcal{F}_{k_{s'}}[s']] \in R^\sharp(k_s) \wedge \\ \text{stop} \in F^\sharp(s) &\iff \text{stop} \in R^\sharp(k_s) \wedge \\ [\text{fn } x \Rightarrow s] \in E^\sharp(y) &\iff [\text{fn } x, k_s \Rightarrow \mathcal{F}_{k_s}[s]] \in R^\sharp(y) \end{aligned}$$

For the purpose of the equivalence we equate the special variables x_r and v_r , both naming the result of the computations. We prove the theorem by combining an implication in each direction with the identity from Lemma 5.1. We formulate both implication as relations and prove that both relations are preserved by the transfer functions.

5.4 ANF-CPS equivalence

We formally define a relation $R_{\text{CPS}}^{\text{ANF}}$ that relates ANF analysis triples to CPS analysis pairs.

Definition 5.2. $\langle C, F^\sharp, E^\sharp \rangle R_{\text{CPS}}^{\text{ANF}} \langle Q^\sharp, R^\sharp \rangle$ iff $\forall s$:

$$\begin{aligned} s \in C &\implies \mathcal{F}_{k_s}[s] \in Q^\sharp \wedge \\ [x, s'] \in F^\sharp(s) &\implies [\text{fn } x \Rightarrow \mathcal{F}_{k_{s'}}[s']] \in R^\sharp(k_s) \wedge \\ \text{stop} \in F^\sharp(s) &\implies \text{stop} \in R^\sharp(k_s) \wedge \\ [\text{fn } x \Rightarrow s] \in E^\sharp(y) &\implies [\text{fn } x, k_s \Rightarrow \mathcal{F}_{k_s}[s]] \in R^\sharp(y) \end{aligned}$$

First we need a small lemma relating the ANF helper function to one of the CPS helper functions.

Lemma 5.2.

$$\begin{aligned} [\text{fn } x \Rightarrow s] \in \mu^\sharp(t, E^\sharp) \wedge \langle C, F^\sharp, E^\sharp \rangle R_{\text{CPS}}^{\text{ANF}} \langle Q^\sharp, R^\sharp \rangle \\ \implies [\text{fn } x, k_s \Rightarrow \mathcal{F}_{k_s}[s]] \in \mu_r^\sharp(\mathcal{V}[t], R^\sharp) \end{aligned}$$

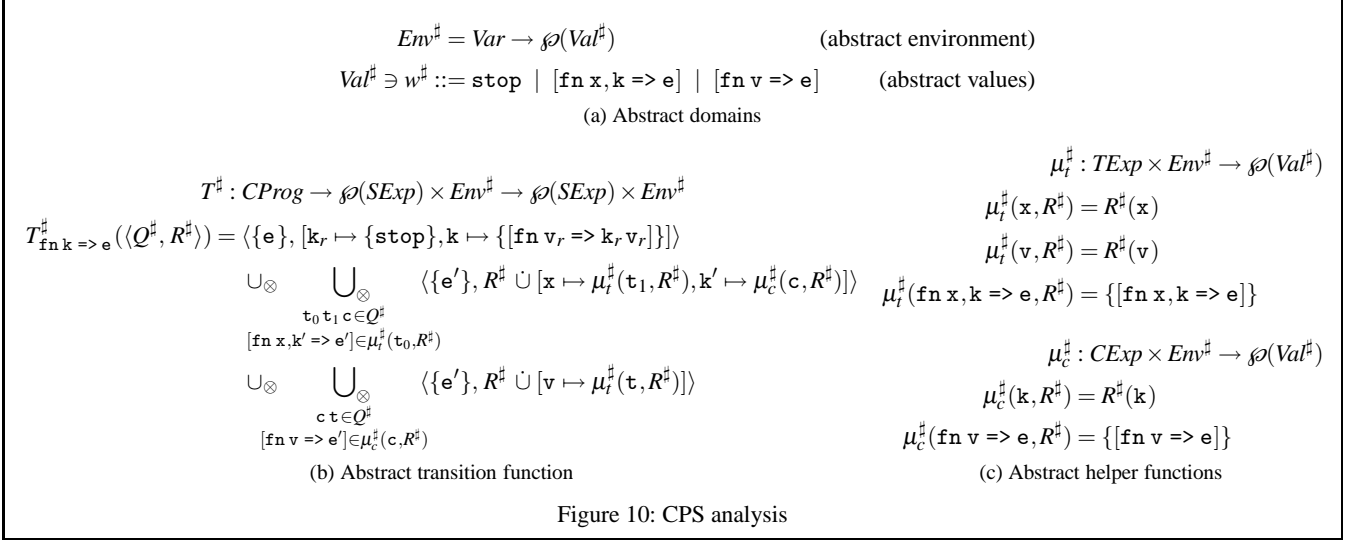


Figure 10: CPS analysis

The relation is preserved by the transfer functions.

Theorem 5.2.

$$\langle C, F^\sharp, E^\sharp \rangle \mathbf{R}_{\text{CPS}}^{\text{ANF}} \langle Q^\sharp, R^\sharp \rangle \implies F_p^\sharp(\langle C, F^\sharp, E^\sharp \rangle) \mathbf{R}_{\text{CPS}}^{\text{ANF}} T_{\mathcal{C}[p]}^\sharp(\langle Q^\sharp, R^\sharp \rangle)$$

Proof. First we name the individual triples of the union in the function body of F^\sharp . We name the first triple of results as initial: $\langle C_I, F_I^\sharp, E_I^\sharp \rangle = \langle \{p\}, [p \mapsto \{[x_r, x_r]\}, x_r \mapsto \{\text{stop}\}], \lambda_{-} \cdot 0 \rangle$. The results of the second, third, fourth, and fifth joined triples corresponding to return, binding, tail call, and non-tail call are named $\langle C_{ret}, F_{ret}^\sharp, E_{ret}^\sharp \rangle$, $\langle C_{bind}, F_{bind}^\sharp, E_{bind}^\sharp \rangle$, $\langle C_{tc}, F_{tc}^\sharp, E_{tc}^\sharp \rangle$ and $\langle C_{ntc}, F_{ntc}^\sharp, E_{ntc}^\sharp \rangle$, respectively. Similarly we name the first result pair in the function body of the CPS analysis as initial: $\langle Q_I^\sharp, R_I^\sharp \rangle = \langle \{e\}, [k_r \mapsto \{\text{stop}\}, k \mapsto \{[\text{fn } v_r \Rightarrow k_r, v_r]\}] \rangle$. The results of the second and third joined pair corresponding to call and return are named $\langle Q_{call}^\sharp, R_{call}^\sharp \rangle$ and $\langle Q_{ret}^\sharp, R_{ret}^\sharp \rangle$, respectively.

The proof proceeds by verifying five relations:

$$\langle C_I, F_I^\sharp, E_I^\sharp \rangle \mathbf{R}_{\text{CPS}}^{\text{ANF}} \langle Q_I^\sharp, R_I^\sharp \rangle \quad (1)$$

$$\langle C_{ret}, F_{ret}^\sharp, E_{ret}^\sharp \rangle \mathbf{R}_{\text{CPS}}^{\text{ANF}} \langle Q_{ret}^\sharp, R_{ret}^\sharp \rangle \quad (2)$$

$$\langle C_{bind}, F_{bind}^\sharp, E_{bind}^\sharp \rangle \mathbf{R}_{\text{CPS}}^{\text{ANF}} \langle Q_{ret}^\sharp, R_{ret}^\sharp \rangle \quad (3)$$

$$\langle C_{tc}, F_{tc}^\sharp, E_{tc}^\sharp \rangle \mathbf{R}_{\text{CPS}}^{\text{ANF}} \langle Q_{call}^\sharp, R_{call}^\sharp \rangle \quad (4)$$

$$\langle C_{ntc}, F_{ntc}^\sharp, E_{ntc}^\sharp \rangle \mathbf{R}_{\text{CPS}}^{\text{ANF}} \langle Q_{call}^\sharp, R_{call}^\sharp \rangle \quad (5)$$

Realizing that the union of related triples and pairs are related we obtain the desired result. \square

After realizing that the bottom elements are related by the above relation, it follows by fixed point induction that their least fixed points (and hence the analyses) are related.

Corollary 5.1. $\text{lfp } F_p^\sharp \mathbf{R}_{\text{CPS}}^{\text{ANF}} \text{lfp } T_{\mathcal{C}[p]}^\sharp$

5.5 CPS-ANF equivalence

Again we formally define a relation now relating CPS analysis pairs to ANF analysis triples.

Definition 5.3. $\langle Q^\sharp, R^\sharp \rangle \mathbf{R}_{\text{ANF}}^{\text{CPS}} \langle C, F^\sharp, E^\sharp \rangle$ iff $\forall e :$

$$e \in Q^\sharp \implies \mathcal{U}[e] \in C \wedge$$

$$[\text{fn } x \Rightarrow e] \in R^\sharp(k_s) \implies [x, \mathcal{U}[e]] \in F^\sharp(s) \wedge$$

$$\text{stop} \in R^\sharp(k_s) \implies \text{stop} \in F^\sharp(s) \wedge$$

$$[\text{fn } x, k_s \Rightarrow e] \in R^\sharp(y) \implies [\text{fn } x \Rightarrow \mathcal{U}[e]] \in E^\sharp(y)$$

We again need a helper lemma relating the helper functions.

Lemma 5.3.

$$[\text{fn } x, k_s \Rightarrow e] \in \mu_t^\sharp(\tau, R^\sharp) \wedge \langle Q^\sharp, R^\sharp \rangle \mathbf{R}_{\text{ANF}}^{\text{CPS}} \langle C, F^\sharp, E^\sharp \rangle$$

$$\implies [\text{fn } x \Rightarrow \mathcal{U}[e]] \in \mu^\sharp(\mathcal{P}[\tau], E^\sharp)$$

This relation is also preserved by the transfer functions.

Theorem 5.3.

$$\langle Q^\sharp, R^\sharp \rangle \mathbf{R}_{\text{ANF}}^{\text{CPS}} \langle C, F^\sharp, E^\sharp \rangle$$

$$\implies T_{\mathcal{C}[p]}^\sharp(\langle Q^\sharp, R^\sharp \rangle) \mathbf{R}_{\text{ANF}}^{\text{CPS}} F_p^\sharp(\langle C, F^\sharp, E^\sharp \rangle)$$

Proof. The proof follows a similar structure to the earlier proof. \square

The bottom elements are related by the relation and it follows by fixed point induction that their least fixed points (and hence the analyses) are related.

Corollary 5.2. $\text{lfp } T_{\mathcal{C}[p]}^\sharp \mathbf{R}_{\text{ANF}}^{\text{CPS}} \text{lfp } F_p^\sharp$

6. Extracting constraints

The resulting analysis may appear complex at first glance. However we can express the analysis in the popular constraint formulation, extracted from the obtained definition. The formulation shown below is in terms of program-specific conditional constraints. Constraints have a (possibly empty) list of preconditions and a conclusion [Palsberg and Schwartzbach, 1995, Gasser et al., 1997]:

$$\{u_1\} \subseteq rhs_1 \wedge \dots \wedge \{u_n\} \subseteq rhs_n \implies lhs \subseteq rhs$$

The constraints operate on the same three domains as the above analysis. Left-hand sides lhs can be of the form $\{u\}$, $F^\sharp([s]_{\equiv})$, or $E^\sharp(x)$, right-hand sides rhs can be of the form C , $F^\sharp([s]_{\equiv})$, or $E^\sharp(x)$, and singleton elements u can be of the form s , c , $[\text{fn } x \Rightarrow s]$, or $[x, s]$. From Fig. 7 we directly read off the following constraints.

i	ANF trace: $\langle C_i, F_i^\sharp, E_i^\sharp \rangle$	CPS trace: $\langle Q_i^\sharp, R_i^\sharp \rangle$
	$\{\text{let } f = \text{fn } x \Rightarrow x \text{ in let } a_1 = f \text{ cn1 in let } a_2 = f \text{ cn2 in } a_2\}$	$\{(\text{fn } f \Rightarrow f \text{ ccn1}(\text{fn } a_1 \Rightarrow f \text{ ccn2}(\text{fn } a_2 \Rightarrow k_p a_2)))(\text{fn } x, k_x \Rightarrow k_x x)\}$
0	$\left[\begin{array}{l} [x_r] \mapsto \{\text{stop}\}, \\ [\text{let } f = \text{fn } x \Rightarrow x \text{ in let } a_1 = f \text{ cn1 in let } a_2 = f \text{ cn2 in } a_2] \mapsto \{[x_r, x_r]\} \end{array} \right]$ $\lambda_{_} . 0$	$\left[\begin{array}{l} k_r \mapsto \{\text{stop}\}, \\ k_p \mapsto \{[\text{fn } v_r \Rightarrow k_r v_r]\} \end{array} \right]$
1	$C_0 \cup \{\text{let } a_1 = f \text{ cn1 in let } a_2 = f \text{ cn2 in } a_2\}$ F_0^\sharp $E_0^\sharp \cup [f \mapsto \{[\text{fn } x \Rightarrow x]\}]$	$Q_0^\sharp \cup \{f \text{ ccn1}(\text{fn } a_1 \Rightarrow f \text{ ccn2}(\text{fn } a_2 \Rightarrow k_p a_2))\}$ $R_0^\sharp \cup [f \mapsto \{[\text{fn } x, k_x \Rightarrow k_x x]\}]$
2	$C_1 \cup \{x\}$ $F_1^\sharp \cup [x \mapsto \{[a_1, \text{let } a_2 = f \text{ cn2 in } a_2]\}]$ $E_1^\sharp \cup [x \mapsto \{\text{cn1}\}]$	$Q_1^\sharp \cup \{k_x x\}$ $R_1^\sharp \cup \left[\begin{array}{l} k_x \mapsto \{[\text{fn } a_1 \Rightarrow f \text{ ccn2}(\text{fn } a_2 \Rightarrow k_p a_2)]\} \\ x \mapsto \{\text{ccn1}\} \end{array} \right]$
3	$C_2 \cup \{\text{let } a_2 = f \text{ cn2 in } a_2\}$ F_2^\sharp $E_2^\sharp \cup [a_1 \mapsto \{\text{cn1}\}]$	$Q_2^\sharp \cup \{f \text{ ccn2}(\text{fn } a_2 \Rightarrow k_p a_2)\}$ $R_2^\sharp \cup [a_1 \mapsto \{\text{ccn1}\}]$
4	C_3 $F_3^\sharp \cup [x \mapsto \{[a_1, \text{let } a_2 = f \text{ cn2 in } a_2], [a_2, a_2]\}]$ $E_3^\sharp \cup [x \mapsto \{\text{cn1}, \text{cn2}\}]$	Q_3^\sharp $R_3^\sharp \cup \left[\begin{array}{l} k_x \mapsto \{[\text{fn } a_1 \Rightarrow f \text{ ccn2}(\text{fn } a_2 \Rightarrow k_p a_2)], [\text{fn } a_2 \Rightarrow k_p a_2]\} \\ x \mapsto \{\text{ccn1}, \text{ccn2}\} \end{array} \right]$
5	$C_4 \cup \{a_2\}$ F_4^\sharp $E_4^\sharp \cup \left[\begin{array}{l} a_1 \mapsto \{\text{cn1}, \text{cn2}\} \\ a_2 \mapsto \{\text{cn1}, \text{cn2}\} \end{array} \right]$	$Q_4^\sharp \cup \{k_p a_2\}$ $R_4^\sharp \cup \left[\begin{array}{l} a_1 \mapsto \{\text{ccn1}, \text{ccn2}\} \\ a_2 \mapsto \{\text{ccn1}, \text{ccn2}\} \end{array} \right]$
6	$C_5 \cup \{x_r\}$ F_5^\sharp $E_5^\sharp \cup [x_r \mapsto \{\text{cn1}, \text{cn2}\}]$	$Q_5^\sharp \cup \{k_r v_r\}$ $R_5^\sharp \cup [v_r \mapsto \{\text{ccn1}, \text{ccn2}\}]$
7	$C_6 \quad F_6^\sharp \quad E_6^\sharp$	$Q_6^\sharp \quad R_6^\sharp$

Table 1: Analysis traces of $\text{let } f = \text{fn } x \Rightarrow x \text{ in let } a_1 = f \text{ cn1 in let } a_2 = f \text{ cn2 in } a_2$ and its CPS transformed counterpart

- For the program p :

$$\{p\} \subseteq C \quad \{[x_r, x_r]\} \subseteq F^\sharp([p]_{\equiv}) \quad \{\text{stop}\} \subseteq F^\sharp([x_r]_{\equiv})$$

- For each return expression t and non-tail call $\text{let } x = t_0 t_1 \text{ in } s'$ in p :

$$\{t\} \subseteq C \wedge \{[x, s']\} \subseteq F^\sharp([t]_{\equiv}) \Rightarrow \begin{cases} \{s'\} \subseteq C \wedge \\ \mu_{\text{sym}}(t, E^\sharp) \subseteq E^\sharp(x) \end{cases}$$

- For each let-binding $\text{let } x = t \text{ in } s$ in p :

$$\{\text{let } x = t \text{ in } s\} \subseteq C \Rightarrow \begin{cases} \{s\} \subseteq C \wedge \\ \mu_{\text{sym}}(t, E^\sharp) \subseteq E^\sharp(x) \end{cases}$$

- For each tail call $t_0 t_1$ and function $\text{fn } x \Rightarrow s'$ in p :

$$\{t_0 t_1\} \subseteq C \wedge \{[\text{fn } x \Rightarrow s']\} \subseteq \mu_{\text{sym}}(t_0, E^\sharp) \Rightarrow \begin{cases} \{s'\} \subseteq C \wedge \\ F^\sharp([t_0 t_1]_{\equiv}) \subseteq F^\sharp([s']_{\equiv}) \wedge \\ \mu_{\text{sym}}(t_1, E^\sharp) \subseteq E^\sharp(x) \end{cases}$$

- For each non-tail call $\text{let } x = t_0 t_1 \text{ in } s$ and function $\text{fn } y \Rightarrow s'$ in p :

$$\{\text{let } x = t_0 t_1 \text{ in } s\} \subseteq C \wedge \{[\text{fn } y \Rightarrow s']\} \subseteq \mu_{\text{sym}}(t_0, E^\sharp) \Rightarrow \begin{cases} \{s'\} \subseteq C \wedge \\ \{[x, s]\} \subseteq F^\sharp([s']_{\equiv}) \wedge \\ \mu_{\text{sym}}(t_1, E^\sharp) \subseteq E^\sharp(y) \end{cases}$$

where we partially evaluate the helper function, i.e., interpret the helper function symbolically at constraint-generation time, to generate a lookup for variables, and a singleton for constants and lambda expressions. The definition of the symbolic helper function otherwise coincides with the abstract helper function μ^\sharp .

We may generate constraints $\{[\text{fn } x \Rightarrow s]\} \subseteq \{[\text{fn } y \Rightarrow s']\}$ of a form not covered by the above grammar. We therefore first preprocess the constraints in linear time, removing vacuously true inclusions $\{[\text{fn } x \Rightarrow s]\} \subseteq \{[\text{fn } x \Rightarrow s]\}$ from each constraint, and removing constraints containing vacuously false preconditions $\{[\text{fn } x \Rightarrow s]\} \subseteq \{w^\sharp\}$, where $[\text{fn } y \Rightarrow s'] \neq w^\sharp$.

The resulting constraint system is formally equivalent to the control flow analysis in the sense that all solutions yield correct control flow information and that the best (smallest) solution of

the constraints is as precise as the information computed by the analysis. More formally:

Theorem 6.1. A solution to the CFA constraints of program p is a safe approximation of the least fixpoint of the analysis function F^{\sharp} induced by p . Furthermore, the least solution to the CFA constraints is equal to the least fixpoint of F^{\sharp} .

Implemented naively, a single constraint may take $O(n)$ space alone. However by using pointers or by labelling each sub-expression and using the pointer/label instead of the sub-expression itself, a single constraint takes only constant space. By linearly determining a representative for each sub-expression, by generating $O(n^2)$ constraints, linear post-processing, and iteratively solving them using a well-known algorithm [Palsberg and Schwartzbach, 1995, Nielson et al., 1999], we can compute the analysis in worst-case $O(n^3)$ time.

The extracted constraints bear similarities to existing constraint-based analyses in the literature. Consider, e.g., calls $t_0 t_1$, which usually gives rise to two conditional constraints [Palsberg, 1995, Nielson et al., 1999]: (1) $\{\{fn\ x \Rightarrow s'\}\} \subseteq \widehat{C}(t_0) \Rightarrow \widehat{C}(t_1) \subseteq \widehat{E}(x)$ and (2) $\{\{fn\ x \Rightarrow s'\}\} \subseteq \widehat{C}(t_0) \Rightarrow \widehat{C}(s') \subseteq \widehat{C}(t_0 t_1)$. The first constraint resembles our third constraint for tail calls. The second “return constraint” differs in that it has a inside-out (or caller-restore) nature, i.e., propagation of return-flow from the function body is handled at the call site. The extracted reachability constraints are similar to Gasser et al. [1997] (modulo an isomorphic encoding $\wp(C) \simeq C \rightarrow \wp(\{\mathbf{on}\})$ of powersets).

7. Conclusion

We have presented a control-flow analysis determining interprocedural control-flow of both calls and returns for a direct-style language. Existing CFAs have focused on analysing which functions are called at a given call site. In contrast, the systematic derivation of our CFA has led to an analysis that provides extra information about where a function returns to at no additional cost. In the presence of tail-call optimization, such information enables the creation of more precise call graphs.

The analysis was developed systematically using Galois connection-based abstract interpretation of a standard operational semantics for that language: the C_aEK abstract machine of Flanagan et al. In addition to being more principled, such a formulation of the analysis is pedagogically pleasing since monomorphism of the analysis is made explicit through two Galois connections: one literally merges all bindings to the same variable and one merges all calling contexts of the same function.

The analysis has been shown to provide a result equivalent to what can be obtained by first CPS transforming the program and then running a control flow analysis derived from a CPS-based operational semantics. This extends previous results obtained by Damian and Danvy, and Palsberg and Wand. The close correspondence between the way that the analyses operate (as illustrated by the analysis trace in Table 1) leads us to conjecture that such equivalence results can be obtained for other CFAs derived using abstract interpretation.

The functional, derived by abstract interpretation, that defines the analysis may appear rather complex at first glance. As a final result, we have shown how to extract from the analysis an equivalent constraint-based formulation expressed in terms of the more familiar conditional constraints. Nevertheless, we stress that the derived functional can be used directly to implement the analysis. We have developed a prototype implementation of the resulting analysis in OCaml.²

² available at <http://www.brics.dk/~jmi/ANF-CFA/>

The analysis has been developed for a minimalistic functional language in order to be able to focus on the abstraction of the control structure induced by function calls and returns. An obvious extension is to enrich the language with numerical operators and study how our Galois connections interact with abstractions such as the interval or polyhedral abstraction of numerical entities.

The calculations involved in the derivation of a CFA are lengthy and would benefit enormously from some form of machine support. *Certified abstract interpretation* [Pichardie, 2005, Cachera et al., 2005] has so far focused on proving the correctness of the analysis inside a proof assistant by using the concretization (γ) component of the Galois connection to prove the correctness of an already defined analysis. Further work should investigate whether proof assistants such as Coq are suitable for conducting the kind of reasoning developed in this paper in a machine-checkable way.

Acknowledgments

The authors thank Matthew Fluet, Amr Sabry, Mitchell Wand, Daniel Damian, Olivier Danvy, and the anonymous referees for comments on earlier versions. Part of this work was done with the support of the Carlsberg Foundation.

References

- J. M. Ashley and R. K. Dybvig. A practical and flexible flow analysis for higher-order languages. *ACM Transactions on Programming Languages and Systems*, 20(4):845–868, 1998.
- A. E. Ayers. Efficient closure analysis with reachability. In M. Billaud, P. Castéran, M.-M. Corsini, K. Musumbu, and A. Rauzy, editors, *Actes WSA'92 Workshop on Static Analysis*, Bigre, pages 126–134, Bordeaux, France, Sept. 1992. Atelier Irisa, IRISA, Campus de Beaulieu.
- D. Cachera, T. Jensen, D. Pichardie, and V. Rusu. Extracting a data flow analyser in constructive logic. *Theoretical Computer Science*, 342(1): 56–78, 2005.
- W. D. Clinger. Proper tail recursion and space efficiency. In K. D. Cooper, editor, *Proc. of the ACM SIGPLAN 1998 Conference on Programming Languages Design and Implementation*, pages 174–185, Montréal, Canada, June 1998.
- C. Consel and O. Danvy. For a better support of static data flow. In J. Hughes, editor, *Proc. of the Fifth ACM Conference on Functional Programming and Computer Architecture*, volume 523 of *LNCS*, pages 496–519, Cambridge, Massachusetts, Aug. 1991. Springer-Verlag.
- P. Cousot. The calculational design of a generic abstract interpreter. In M. Broy and R. Steinbrüggen, editors, *Calculational System Design*. NATO ASI Series F. IOS Press, Amsterdam, 1999.
- P. Cousot. Semantic foundations of program analysis. In S. S. Muchnick and N. D. Jones, editors, *Program Flow Analysis: Theory and Applications*, chapter 10, pages 303–342. Prentice-Hall, 1981.
- P. Cousot and R. Cousot. Abstract interpretation of algebraic polynomial systems. In M. Johnson, editor, *Proc. of the Sixth International Conference on Algebraic Methodology and Software Technology, AMAST'97*, volume 1349 of *LNCS*, pages 138–154, Sydney, Australia, Dec. 1997. Springer-Verlag.
- P. Cousot and R. Cousot. Higher-order abstract interpretation (and application to compartment analysis generalizing strictness, termination, projection and PER analysis of functional languages), invited paper. In H. Bal, editor, *Proc. of the Fifth IEEE International Conference on Computer Languages*, pages 95–112, Toulouse, France, May 1994.
- P. Cousot and R. Cousot. Abstract interpretation frameworks. *Journal of Logic and Computation*, 2(4):511–547, Aug. 1992a.
- P. Cousot and R. Cousot. Abstract interpretation and application to logic programs. *Journal of Logic Programming*, 13(2–3):103–179, 1992b.
- P. Cousot and R. Cousot. Systematic design of program analysis frameworks. In B. K. Rosen, editor, *Proc. of the Sixth Annual ACM Sym-*

- posium on Principles of Programming Languages*, pages 269–282, San Antonio, Texas, Jan. 1979.
- D. Damian and O. Danvy. Syntactic accidents in program analysis: On the impact of the CPS transformation. *Journal of Functional Programming*, 13(5):867–904, 2003. A preliminary version was presented at the 2000 ACM SIGPLAN International Conference on Functional Programming.
- O. Danvy. Three steps for the CPS transformation. Technical Report CIS-92-2, Kansas State University, Manhattan, Kansas, Dec. 1991.
- B. A. Davey and H. A. Priestley. *Introduction to Lattices and Order*. Cambridge University Press, Cambridge, England, second edition, 2002.
- S. K. Debray and T. A. Proebsting. Interprocedural control flow analysis of first-order programs with tail-call optimization. *ACM Transactions on Programming Languages and Systems*, 19(4):568–585, 1997.
- C. Flanagan, A. Sabry, B. F. Duba, and M. Felleisen. The essence of compiling with continuations. In D. W. Wall, editor, *Proc. of the ACM SIGPLAN 1993 Conference on Programming Languages Design and Implementation*, pages 237–247, Albuquerque, New Mexico, June 1993.
- K. L. S. Gasser, F. Nielson, and H. R. Nielson. Systematic realisation of control flow analyses for CML. In M. Tofte, editor, *Proc. of the Second ACM SIGPLAN International Conference on Functional Programming*, pages 38–51, Amsterdam, The Netherlands, June 1997.
- R. Giacobazzi, F. Ranzato, and F. Scozzari. Making abstract interpretations complete. *J. ACM*, 47(2):361–416, 2000.
- N. Heintze. Set-based program analysis of ML programs. In C. L. Talcott, editor, *Proc. of the 1994 ACM Conference on Lisp and Functional Programming*, LISP Pointers, Vol. VII, No. 3, pages 306–317, Orlando, Florida, June 1994.
- N. D. Jones. Flow analysis of lambda expressions (preliminary version). In S. Even and O. Kariv, editors, *Automata, Languages and Programming, 8th Colloquium, Acre (Akko)*, volume 115 of LNCS, pages 114–128, Israel, July 1981. Springer-Verlag.
- J. Midtgaard. Control-flow analysis of functional programs. Technical Report BRICS RS-07-18, Dept. of Comp. Sci., University of Aarhus, Aarhus, Denmark, Dec. 2007. Accepted for publication in *ACM Computing Surveys*.
- J. Midtgaard and T. Jensen. A calculational approach to control-flow analysis by abstract interpretation. In M. Alpuente and G. Vidal, editors, *Static Analysis, 15th International Symposium, SAS 2008*, volume 5079 of LNCS, pages 347–362, Valencia, Spain, July 2008a. Springer-Verlag.
- J. Midtgaard and T. P. Jensen. Control-flow analysis of function calls and returns by abstract interpretation. Rapport de Recherche RR-6681, INRIA Rennes – Bretagne Atlantique, Oct. 2008b.
- M. Might and O. Shivers. Environmental analysis via ΔCFA. In S. Peyton Jones, editor, *Proc. of the 33rd Annual ACM Symposium on Principles of Programming Languages*, pages 127–140, Charleston, South Carolina, Jan. 2006.
- R. Milner and M. Tofte. Co-induction in relational semantics. *Theoretical Computer Science*, 87(1):209–220, 1991.
- F. Nielson and H. R. Nielson. Infinitary control flow analysis: a collecting semantics for closure analysis. In N. D. Jones, editor, *Proc. of the 24th Annual ACM Symposium on Principles of Programming Languages*, pages 332–345, Paris, France, Jan. 1997.
- F. Nielson, H. R. Nielson, and C. Hankin. *Principles of Program Analysis*. Springer-Verlag, 1999.
- H. R. Nielson and F. Nielson. Flow logic: a multi-paradigmatic approach to static analysis. In T. Æ. Mogensen, D. A. Schmidt, and I. H. Sudborough, editors, *The Essence of Computation: Complexity, Analysis, Transformation. Essays Dedicated to Neil D. Jones*, volume 2566 of LNCS, pages 223–244. Springer-Verlag, 2002.
- J. Palsberg. Closure analysis in constraint form. *ACM Transactions on Programming Languages and Systems*, 17(1):47–62, 1995.
- J. Palsberg and M. I. Schwartzbach. Safety analysis versus type inference. *Information and Computation*, 118(1):128–141, 1995.
- J. Palsberg and M. Wand. CPS transformation of flow information. *Journal of Functional Programming*, 13(5):905–923, 2003.
- D. Pichardie. *Interprétation abstraite en logique intuitioniste: extraction d'analyseurs Java certifiés*. PhD thesis, Université de Rennes 1, Sept. 2005.
- A. Sabry and M. Felleisen. Is continuation-passing useful for data flow analysis? In V. Sarkar, editor, *Proc. of the ACM SIGPLAN 1994 Conference on Programming Languages Design and Implementation*, pages 1–12, Orlando, Florida, June 1994.
- P. Sestoft. Replacing function parameters by global variables. In J. E. Stoy, editor, *Proc. of the Fourth International Conference on Functional Programming and Computer Architecture*, pages 39–53, London, England, Sept. 1989.
- O. Shivers. Control-flow analysis in Scheme. In M. D. Schwartz, editor, *Proc. of the ACM SIGPLAN 1988 Conference on Programming Languages Design and Implementation*, pages 164–174, Atlanta, Georgia, June 1988.
- F. Spoto and T. P. Jensen. Class analyses as abstract interpretations of trace semantics. *ACM Transactions on Programming Languages and Systems*, 25(5):578–630, 2003.

A. Underlying mathematical material

This section is based on known material [Cousot and Cousot, 1979, Cousot, 1981, Cousot and Cousot, 1992b, 1994, Davey and Priestley, 2002].

A complete lattice is a partially ordered set $\langle C; \sqsubseteq, \perp, \top, \sqcup, \sqcap \rangle$ (poset), such that the least upper bound $\sqcup S$ and the greatest lower bound $\sqcap S$ exists for every subset S of C . $\perp = \sqcap C$ denotes the infimum of C and $\top = \sqcup C$ denotes the supremum of C . The set of total functions $D \rightarrow C$, whose domain is a complete lattice $\langle C; \sqsubseteq, \perp, \top, \sqcup, \sqcap \rangle$, is itself a complete lattice $\langle D \rightarrow C; \sqsubseteq, \perp, \top, \sqcup, \sqcap \rangle$ under the pointwise ordering $f \sqsubseteq f' \iff \forall x. f(x) \sqsubseteq f'(x)$, with bottom, top, join, and meet extended similarly. The powersets $\wp(S)$ of a set S ordered by set inclusion is a complete lattice $\langle \wp(S); \subseteq, \emptyset, S, \cup, \cap \rangle$.

A Galois connection is a pair of functions α, γ between two posets $\langle C; \sqsubseteq \rangle$ and $\langle A; \leq \rangle$ such that for all $a \in A, c \in C$: $\alpha(c) \leq a \iff c \sqsubseteq \gamma(a)$. Equivalently a Galois connection can be defined as a pair of functions satisfying (a) α and γ are monotone, (b) $\alpha \circ \gamma$ is reductive, and (c) $\gamma \circ \alpha$ is extensive. Galois connections are typeset $\langle C; \sqsubseteq \rangle \xrightarrow[\alpha]{\gamma} \langle A; \leq \rangle$. We omit the orderings when they are clear from the context. For a Galois connection between two complete lattices α is a complete join-morphism (CJM) and γ is a complete meet morphism. The composition of two Galois connections $\langle C; \sqsubseteq \rangle \xrightarrow[\alpha_1]{\gamma_1} \langle B; \sqsubseteq \rangle$ and $\langle B; \sqsubseteq \rangle \xrightarrow[\alpha_2]{\gamma_2} \langle A; \leq \rangle$ is itself a Galois connection $\langle C; \sqsubseteq \rangle \xrightarrow[\alpha_2 \circ \alpha_1]{\gamma_1 \circ \gamma_2} \langle A; \leq \rangle$. Galois connections in which α is surjective (or equivalently γ is injective) are typeset $\langle C; \sqsubseteq \rangle \xrightarrow[\alpha]{\gamma} \langle A; \leq \rangle$. Galois connections in which γ is surjective (or equivalently α is injective) are typeset $\langle C; \sqsubseteq \rangle \xleftarrow[\alpha]{\gamma} \langle A; \leq \rangle$. When both α and γ are surjective, the two domains are isomorphic.

A(n upper) closure operator ρ is map $\rho : S \rightarrow S$ on a poset $\langle S; \sqsubseteq \rangle$, that is (a) monotone: (for all $s, s' \in S$: $s \sqsubseteq s' \implies \rho(s) \sqsubseteq \rho(s')$), (b) extensive: (for all $s \in S$: $s \sqsubseteq \rho(s)$), and (c) idempotent, (for all $s \in S$: $\rho(s) = \rho(\rho(s))$). A closure operator ρ induces a Galois connection $\langle S; \sqsubseteq \rangle \xrightarrow[\rho]{1} \langle \rho(S); \sqsubseteq \rangle$, writing $\rho(S)$ for $\{\rho(s) \mid s \in S\}$ and 1 for the identity function. Furthermore the image of a complete lattice $\langle C; \sqsubseteq, \perp, \top, \sqcup, \sqcap \rangle$ by an upper closure operator is itself a complete lattice $\langle \rho(C); \sqsubseteq, \rho(\perp), \top, \lambda X. \rho(\sqcup X), \sqcap \rangle$.