A Verified Compiler from Isabelle/HOL to CakeML

LARS HUPEL, Technische Universität München
TOBIAS NIPKOW, Technische Universität München

Many theorem provers can generate functional programs from definitions or proofs. However, this code generation needs to be trusted. Except for the HOL4 system, which has a proof producing code generator for a subset of ML. We go one step further and provide a verified compiler from Isabelle/HOL to CakeML. More precisely we combine a simple proof producing translation of recursion equations in Isabelle/HOL into a deeply embedded term language with a fully verified compilation chain to the target language CakeML.

CCS Concepts: •**Theory of computation → Logic and verification; Program semantics;** *Higher order logic;* •**Software and its engineering → Compilers;**

Additional Key Words and Phrases: Isabelle, CakeML, compiler, higher-order term rewriting

ACM Reference format:
Lars Hupel and Tobias Nipkow. 2017. A Verified Compiler from Isabelle/HOL to CakeML. *PACM Progr. Lang.* 1, 1, Article 1 (January 2017), 28 pages.
DOI: 10.1145/nnnnnnn.nnnnnnn

1 INTRODUCTION

Many theorem provers have the ability to generate executable code in some (typically functional) programming language from definitions, lemmas and proofs (e.g. [7, 9, 10, 12, 16, 26, 36]). This makes code generation part of the trusted kernel of the system. Myreen and Owens [29] closed this gap for the HOL4 system: they have implemented a tool that translates from HOL4 into *CakeML*, a subset of SML, and proves a theorem stating that a result produced by the CakeML code is correct w.r.t. the HOL functions. They also have a verified implementation of CakeML [23, 38]. We go one step further and provide a once-and-for-all verified compiler from (deeply embedded) function definitions in Isabelle/HOL [31, 32] into CakeML proving partial correctness of the generated CakeML code w.r.t. the original functions. This is like the step from dynamic to static type checking. It also means that preconditions on the input to the compiler are explicitly given in the correctness theorem rather than implicitly by a failing translation. To the best of our knowledge this is the first verified (as opposed to certifying) compiler from function definitions in a logic into a programming language.

Our compiler is composed of multiple phases and in principle applicable to other languages than Isabelle/HOL or even HOL:

- We erase types right away. Hence the type system of the source language is irrelevant.
- We merely assume that the source language has a semantics based on equational logic.

The compiler operates in three stages:

(1) The preprocessing phase eliminates features that are not supported by our compiler. Most importantly, *dictionary construction* eliminates occurrences of type classes in HOL terms. It introduces dictionary datatypes and new constants and proves the equivalence of old and new constants (§7).

2017. 2475-1421/2017/1-ART1 $15.00
DOI: 10.1145/nnnnnnn.nnnnnnn

(2) The *deep embedding* lifts HOL terms into terms of type term, a HOL model of HOL terms. For each constant c (of arbitrary type) it defines a constant c' of type term and proves a theorem that expresses equivalence (§3).
(3) There are multiple *compiler phases* that eliminate certain constructs from the term type, until we arrive at the CakeML expression type. Most phases target a different intermediate term type (§5).

The first two stages are preprocessing, are implemented in ML and produce certificate theorems. Only these stages are specific to Isabelle. The third (and main) stage is implemented completely in the logic HOL, without recourse to ML. Its correctness is verified once and for all. All Isabelle definitions and proofs can be found in the supplementary material.

2 RELATED WORK

There is existing work in the Coq [2, 15] and HOL [29] communities for proof producing or verified extraction of functions defined in the logic into executable. Anand *et al.* [2] present work in progress on a verified compiler from Gallina (Coq's specification language) via untyped intermediate languages to CompCert C light. They plan to connect their extraction routine to the CompCert compiler [25].

Translation of type classes into dictionaries is an important feature of Haskell compilers. In the setting of Isabelle/HOL, this has been described by Wenzel [42] and Krauss *et al.* [22]. Haftmann and Nipkow [17] use this construction to compile HOL definitions into target languages that do not support type classes, e.g. Standard ML and OCaml. In this work, we provide a certifying translation that eliminates type classes inside the logic.

Compilation of pattern matching is well understood in literature [3, 34, 37]. In this work, we contribute a transformation of sets of equations with pattern matching on the left-hand side into a single equation with nested pattern matching on the right-hand side. This is implemented and verified inside Isabelle.

Besides CakeML, there are many projects for verified compilers for functional programming languages of various degrees of sophistication and realism (e.g. [5, 11, 14]). Particularly modular is the work by Neis *et al.* [30] on a verified compiler for an ML-like imperative source language. The main distinguishing feature of our work is that we start from a set of higher-order recursion equations with pattern matching on the left-hand side rather than a lambda calculus with pattern matching on the right-hand side. On the other hand we stand on the shoulders of CakeML which allows us to bypass all complications of machine code generation. Note that much of our compiler is not specific to CakeML and that it would be possible to retarget it to, for example, Pilsner abstract syntax with moderate effort.

3 DEEP EMBEDDING

Starting with a HOL definition, we derive a new, *reified* definition in a deeply embedded term language depicted in Figure 1a. This term language corresponds closely to the term datatype of Isabelle's implementation (using de Bruijn indices [13]), but without types and schematic variables.

To establish a formal connection between the original and the reified definitions, a "family of relations" is used, "defined by induction on types".[1] This concept of a *logical relation* is well-understood in literature [20] and can be nicely implemented in Isabelle using type classes. Note that the use of type classes here is restricted to correctness proofs; it is not required for the execution of the compiler itself. That way, there is no contradiction to the elimination of type classes occurring in a previous stage.

[1] https://ncatlab.org/nlab/revision/logical+relation/7

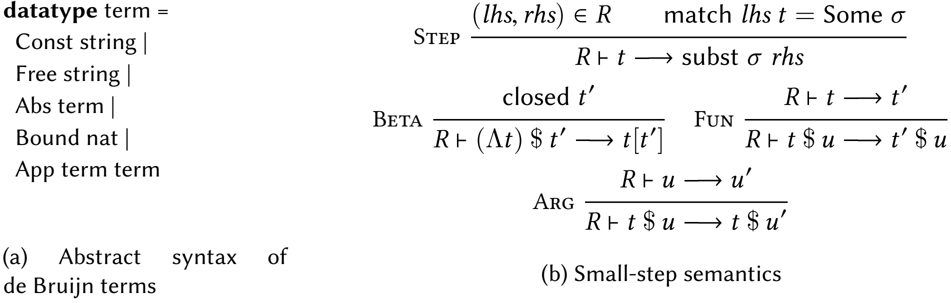

(a) Abstract syntax of de Bruijn terms

(b) Small-step semantics

Fig. 1. Basic syntax and semantics of the term type

Notation. We abbreviate App $t\ u$ to $t\ \$\ u$ and Abs t to $\Lambda\ t$. Other term types introduced later in this paper use the same conventions. We reserve λ for abstractions in HOL itself. Typing judgments are written with a double colon: $t :: \tau$.

Embedding operation. Embedding is implemented in ML. We denote this operation using angle brackets: $\langle t \rangle$, where t is an arbitrary HOL expression and the result $\langle t \rangle$ is a HOL value of type term. It is a purely syntactic transformation, without preliminary evaluation or reduction, and it discards type information. The following examples illustrate this operation and typographical conventions concerning variables and constants:

$$\langle x \rangle = \mathsf{Free}\ \text{"x"}$$
$$\langle \mathsf{f} \rangle = \mathsf{Const}\ \text{"f"}$$
$$\langle \lambda x\ y.\ \mathsf{f}\ y\ x \rangle = \Lambda\ (\Lambda\ (\langle \mathsf{f} \rangle\ \$\ \mathsf{Bound}\ 0\ \$\ \mathsf{Bound}\ 1))$$

Small-step semantics. Figure 1b specifies the small-step semantics for term. It is reminiscent of *higher-order term rewriting*, and modelled closely after equality in HOL. The basic idea is that if the proposition $t = u$ can be proved equationally in HOL (without symmetry), then $R \vdash \langle t \rangle \longrightarrow^* \langle u \rangle$ holds (where $R ::$ (term $\times$ term) set). We call R the *rule set.* It is the result of translating a set of defining equations $lhs = rhs$ into pairs $(\langle lhs \rangle, \langle rhs \rangle) \in R$.

Rule Step performs a rewrite step by picking a rewrite rule from R and rewriting the term at the root. For that purpose, match and subst are (mostly) standard first-order matching and substitution (see §4 for details).

Rule Beta performs β-reduction. Type term represents bound variables by de Bruijn indices. The notation $t[t']$ represents the substitution of the outermost bound variable in t with t'.

Our semantics does not constitute a fully-general higher-order term rewriting system, because we do not allow substitution under binders. For de Bruijn terms, this would pose no problem, but as soon as we introduce named bound variables, substitution under binders requires dealing with capture. To avoid this altogether, all our semantics expect terms that are substituted into abstractions to be closed. However, this does not mean that we restrict ourselves to any particular evaluation order. Both call-by-value and call-by-name can be used in the small-step semantics. But later on, the target semantics will only use call-by-value.

Embedding relation. We denote the concept that an embedded term t corresponds to a HOL term a of type τ w.r.t. rule set R with the syntax $R \vdash t \approx a$. If we want to be explicit about the type, we index the relation: $\approx_\tau$.

datatype nterm =	**datatype** pterm =	**datatype** sterm =
Nconst string | Nvar string |	Pconst string | Pvar string |	Sconst string | Svar string |
Nabs term |	Pabs ((term × pterm) set) |	Sabs ((term × sterm) list) |
Napp term term	Papp pterm pterm	Sapp sterm sterm
(a) Named bound variable	(b) Explicit pattern matching	(c) Sequential clauses

Fig. 2. Intermediate term types

For ground types, this can be defined easily. For example, the following two rules define $\approx_{\text{nat}}$:

$$\frac{}{R \vdash \langle 0 \rangle \approx_{\text{nat}} 0} \qquad \frac{R \vdash \langle t \rangle \approx_{\text{nat}} n}{R \vdash \langle \text{Suc } t \rangle \approx_{\text{nat}} \text{Suc } n}$$

Definitions of $\approx$ for arbitrary datatypes without nested recursion can be derived mechanically in the same fashion as for nat, where they constitute one-to-one relations. Note that for ground types, $\approx$ ignores R. The reason why $\approx$ is parametrized on R will become clear in a moment.

For function types, we follow Myreen and Owen's approach [29]. The statement $R \vdash t \approx f$ can be interpreted as "$t \mathbin{\$} \langle a \rangle$ can be rewritten to $\langle f\ a \rangle$ for all a". Because this might involve applying a function definition from R, the $\approx$ relation must be indexed by the rule set. As a notational convenience, we define another relation $R \vdash t \downarrow x$ to mean that there is a t' such that $R \vdash t \longrightarrow^* t'$ and $R \vdash t' \approx x$. Using this notation, we formally define $\approx$ for functions as follows:

$$R \vdash t \approx f \leftrightarrow (\forall x\ t_x.\ R \vdash t_x \downarrow x \longrightarrow R \vdash t \mathbin{\$} t_x \downarrow f\ x)$$

Example. As a running example, we will use the map function on lists:

$$\begin{array}{llll} \text{map} & f & [] & = & [] \\ \text{map} & f & (x \# xs) & = & f\ x \# \text{map}\ f\ xs \end{array}$$

The result of embedding this function is a set of rules map′:

```
map' =
 {(Const "List.list.map" $ Free "f" $ (Const "List.list.Cons" $ Free "x21" $ Free "x22"),
    Const "List.list.Cons" $ (Free "f" $ Free "x21") $ (Const "List.list.map" $ Free "f" $ Free "x22")),
  (Const "List.list.map" $ Free "f" $ Const "List.list.Nil",
    Const "List.list.Nil")|}
```

together with the theorem map′ ⊢ Const "List.list.map" $\downarrow$ map, which is proven by simple induction over map. Constant names like "List.list.map" come from the fully-qualified internal names in HOL.

The induction principle for the proof arises from the use of the **fun** command that is used to define recursive functions in HOL [21]. But the user is also allowed to specify custom equations for functions, in which case we will use heuristics to generate and prove the appropriate induction theorem. For simplicity, we will use the term *(defining) equation* uniformly to refer to any set of equations, either default ones or ones specified by the user.

4 TERMS, MATCHING AND SUBSTITUTION

The compiler transforms the initial term type through various stages. All intermediate types are depicted in Figure 1a and 2. This section gives an overview and introduces necessary terminology.

Preliminaries. The function arrow in HOL is $\Rightarrow$. The cons operator on lists is the infix #.

Throughout the paper, the concept of *mappings* is pervasive: We use the type notation $\alpha \rightharpoonup \beta$ to denote a function $\alpha \Rightarrow \beta$ option. In certain contexts, a mapping may also be called an *environment*. We write mapping literals using brackets: $[a \Rightarrow x, b \Rightarrow y, \ldots]$. If it is clear from the context that σ is defined on a, we often treat the lookup $\sigma\ a$ as returning an $x :: \beta$.

The functions dom $:: (\alpha \rightharpoonup \beta) \Rightarrow \alpha$ set and range $:: (\alpha \rightharpoonup \beta) \Rightarrow \beta$ set return the *domain* and *range* of a mapping, respectively.

Dropping entries from a mapping is denoted by $\sigma - k$, where σ is a mapping and k is either a single key or a set of keys. We use $\sigma' \subseteq \sigma$ to denote that σ' is a sub-mapping of σ, that is, dom $\sigma' \subseteq$ dom σ and $\forall a \in \text{dom}\ \sigma'.\ \sigma'\ a = \sigma\ a$.

Merging two mappings σ and ρ is denoted with $\sigma +\!\!+ \rho$. It constructs a new mapping with the union domain of σ and ρ. Entries from ρ override entries from σ. That is, $\rho \subseteq \sigma +\!\!+ \rho$ holds, but not necessarily $\sigma \subseteq \sigma +\!\!+ \rho$.

All mappings and sets are assumed to be finite. In the formalization, this is enforced by using subtypes of $\rightharpoonup$ and set. Note that one cannot define datatypes by recursion through sets for cardinality reasons. However, for finite sets, it is possible. This is required to construct the various term types. We leverage facilities of Blanchette *et al.*'s **datatype** command to define these subtypes [8].

Standard functions. All type constructors that we use ($\rightharpoonup$, set, list, option, ...) support the standard operations map and rel. For lists, map is the regular covariant map. For mappings, the function has the type $(\beta \Rightarrow \gamma) \Rightarrow (\alpha \rightharpoonup \beta) \Rightarrow (\alpha \rightharpoonup \gamma)$. It leaves the domain unchanged, but applies a function to the range of the mapping.

Function rel_τ lifts a binary predicate $P :: \alpha \Rightarrow \alpha \Rightarrow$ bool to the type constructor τ. We call this lifted relation the *relator* for a particular type.

For datatypes, its definition is structural, for example:

$$\frac{}{\text{rel}_{\text{list}}\ P\ []\ []} \qquad \frac{\text{rel}_{\text{list}}\ P\ xs\ ys \qquad P\ x\ y}{\text{rel}_{\text{list}}\ P\ (x \# xs)\ (y \# ys)}$$

For sets and mappings, the definition is a little bit more subtle.

Definition 4.1 (Set relator). For each element $a \in A$, there must be a corresponding element $b \in B$ such that $P\ a\ b$, and vice versa. Formally:

$$\text{rel}_{\text{set}}\ P\ A\ B \leftrightarrow (\forall x \in A.\ \exists y \in B.\ P\ x\ y) \wedge (\forall y \in B.\ \exists x \in A.\ P\ x\ y)$$

Definition 4.2 (Mapping relator). For each a, $m\ a$ and $n\ a$ must be related according to $\text{rel}_{\text{option}}\ P$. Formally:

$$\text{rel}_{\text{mapping}}\ P\ m\ n \leftrightarrow (\forall a.\ \text{rel}_{\text{option}}\ P\ (m\ a)\ (n\ a))$$

Term types. There are four distinct term types: term, nterm, pterm, and sterm. All of them support the notions of free variables, matching and substitution. Free variables are always a finite set of strings. Matching a term against a *pattern* yields an optional mapping of type string $\rightharpoonup \alpha$ from free variable names to terms.

Note that the type of patterns is itself term instead of a dedicated pattern type. The reason is that we have to subject patterns to a linearity constraint anyway and may use this constraint to carve out the relevant subset of terms:

Definition 4.3. A term is *linear* if there is at most one occurrence of any variable, it contains no abstractions, and in an application $f \mathbin{\$} x$, f must not be a free variable. The HOL predicate is called linear :: term $\Rightarrow$ bool.

Because of the similarity of operations across the term types, they are all instances of the term type class. Note that in Isabelle, classes and types live in different namespaces. The term type and the term type class are separate entities.

Definition 4.4. A *term type* τ supports the operations match :: term $\Rightarrow \tau \Rightarrow$ (string $\rightharpoonup \tau$), subst :: (string $\rightharpoonup \tau) \Rightarrow \tau \Rightarrow \tau$ and frees :: $\tau \Rightarrow$ string set together with the following axioms:

$$\begin{array}{rcrl} & & \text{subst } [] \; t & = t \\ x \notin \text{frees } t & \longrightarrow & \text{subst } (\sigma - x) \; t & = \text{subst } \sigma \; t \\ \text{match } p \; t = \text{Some } \sigma & \longrightarrow & \text{dom } \sigma & = \text{frees } p \\ \text{match } p \; t = \text{Some } \sigma & \longrightarrow & \forall u \in \text{range } \sigma.\ \text{frees } u & \subseteq \text{frees } t \\ \forall u \in \text{range } \sigma.\ \text{frees } u = \emptyset & \longrightarrow & \text{frees } (\text{subst } \sigma \; t) & = \text{frees } t - \text{dom } \sigma \end{array}$$

We also define the following derived functions:

- matchs matches a list of patterns and terms sequentially, producing a single mapping
- closed t is an abbreviation for frees $t = \emptyset$
- closed σ is an overloading of closed, denoting that all values in a mapping are closed

The above set of axioms does not strive to fully specify an abstract term algebra (cf. work by Schmidt-Schauß and Siekmann [35]). Instead, they are chosen according to the needs of this formalization.

Definition 4.5. An *equation* is a pair of a pattern *(left-hand side)* and a term *(right-hand side)*. The pattern is of the form $f \mathbin{\$} p_1 \mathbin{\$} \ldots \mathbin{\$} p_n$, where f is a constant (i.e. of the form Const *name*). We refer to both f or *name* interchangeably as the *function symbol* of the equation.

Following term rewriting terminology, we sometimes refer to an equation as *rule*.

4.1 De Bruijn terms (term)

The definition of term is almost an exact copy of Isabelle's internal term type, with the notable omissions of type information and schematic variables (Figure 1a). The implementation of β-reduction is straightforward via index shifting of bound variables.

Matching works largely the same across the different term types:

```
fun match :: term ⇒ term ⇒ (string ⇀ term) option where
match (Const x) (Const y) = (if x = y then Some [] else None)
match (t1 $ t2) (u1 $ u2) = do {
  env1 ← match t1 u1;
  env2 ← match t2 u2;
  Some (env1 ++ env2)
}
match (Free s) t = Some [s ↦ t]
match _ _ = None
```

The first argument denotes the pattern, the second the object. A notable deviation from matching as discussed in term rewriting literature is that the result of matching is only well-defined if the pattern is linear.

4.2 Named bound variables (nterm)

The nterm type is similar to term, but removes the distinction between *bound* and *free* variables (Figure 2a). Instead, there are only named variables. As mentioned in the previous section, we forbid substitution of terms that are not closed in order to avoid capture. This is also reflected in the syntactic side conditions of the correctness proofs (§5.1).

4.3 Explicit pattern matching (pterm)

Functions in HOL are usually defined using *implicit* pattern matching, that is, the terms p_i occurring on the left-hand side $\langle \mathsf{f}\ p_1\ \ldots\ p_n \rangle$ of an equation must be constructor patterns. This is also common among functional programming languages like Haskell or OCaml. CakeML only supports *explicit* pattern matching using case expressions. A function definition consisting of multiple defining equations must hence be translated to the form $f = \lambda x.\ \textbf{case}\ x\ \textbf{of}\ \ldots$. The elimination proceeds by iteratively removing the last parameter in the block of equations until none are left.

In our formalization, we opted to combine the notion of abstraction and case expression, yielding *case abstractions*, represented as the Pabs constructor in Figure 2b. This is similar to the fn construct in Standard ML, which denotes an anonymous function that immediately matches on its argument [27]. The same construct also exists in Haskell with the LambdaCase language extension.[2] We chose this representation mainly for two reasons: First, it allows for a simpler language grammar because there is only one (shared) constructor for abstraction and case expression. Second, the elimination procedure outlined above does not have to introduce fresh names in the process. Later, when translating to CakeML syntax, fresh names are introduced and proved correct in a separate step.

The set of pairs of pattern and right-hand side inside a case abstraction is referred to as *clauses.* As a short-hand notation, we use $\Lambda\{p_1 \Rightarrow t_1, p_2 \Rightarrow t_2, \ldots\}$.

4.4 Sequential clauses (sterm)

In the term rewriting fragment of HOL, the order of rules is not significant. If a rule matches, it can be applied, regardless when it was defined or proven. This is reflected by the use of sets in the rule and term types. For CakeML, the rules need to be applied in a deterministic order, i.e. sequentially. The sterm type only differs from pterm by using list instead of set (Figure 2c). Hence, case abstractions use list brackets: $\Lambda[p_1 \Rightarrow t_1, p_2 \Rightarrow t_2, \ldots]$.

4.5 Irreducible terms (value)

CakeML distinguishes between *expressions* and *values.* Whereas expressions may contain free variables or β-redexes, values are closed and fully evaluated. Both have a notion of abstraction, but values differ from expressions in that they contain an environment binding free variables.

Consider the expression $(\lambda x.\lambda y.x)\ (\lambda z.z)$, which is rewritten (by β-reduction) to $\lambda y.\lambda z.z$. Note how the bound variable x disappears, since it is replaced. This is contrary to how programming languages are usually implemented: evaluation does not happen by substituting the argument term t for the bound variable x, but by recording the binding $x \mapsto t$ in an environment [23]. A pair of an abstraction and an environment is usually called a *closure* [24, 39].

In CakeML, this means that evaluation of the above expression results in the closure

$$(\lambda y.x, [\text{"x"} \mapsto (\lambda z.z, [])])$$

Note the nested structure of the closure, whose environment itself contains a closure.

To reflect this in our formalization, we introduce a type value of values (explanation inline):

```
datatype value =
  (* constructor value: a data constructor applied to multiple values *)
  Vconstr string (value list) |
  (* closure: clauses combined with an environment mapping variables to values *)
  Vabs ((term × sterm) list) (string ⇀ value) |
  (* recursive closures: a group of mutually recursive function bodies with an environment mapping variables to values *)
```

[2] https://downloads.haskell.org/~ghc/7.8.4/docs/html/users_guide/syntax-extns.html

$\mathsf{Vrecabs}\ (\mathsf{string} \rightharpoonup ((\mathsf{term} \times \mathsf{sterm})\ \mathsf{list}))\ \mathsf{string}\ (\mathsf{string} \rightharpoonup \mathsf{value})$

The above example evaluates to the closure:

$$\mathsf{Vabs}\ \big[\ \langle y\rangle \Rightarrow \langle x\rangle\ \big]\ \big[\text{"x"} \mapsto \mathsf{Vabs}\ [\langle z\rangle \Rightarrow \langle z\rangle]\ []\big]$$

The third case for recursive closures only becomes relevant when we conflate variables and constants. As long as the rule set *rs* is kept separate, recursive calls are straightforward: the appropriate definition for the constant can be looked up there. CakeML knows no such distinction between constants and variables, hence everything has to reside in a single environment σ.

Consider this example of odd and even:

$$\mathsf{odd}\ 0 = \mathsf{False} \qquad \mathsf{even}\ 0 = \mathsf{True}$$
$$\mathsf{odd}\ (\mathsf{Suc}\ n) = \mathsf{even}\ n \qquad \mathsf{even}\ (\mathsf{Suc}\ n) = \mathsf{odd}\ n$$

When evaluating the term odd k, the definitions of even and odd themselves must be available in the environment captured in the definition of odd. However, it would be cumbersome in HOL to construct such a Vabs that refers to itself. Instead, we capture the expressions used to define odd and even in a recursive closure. Other encodings might be possible, but since we are targeting CakeML, we are opting to model it in a similar way as its authors do.

For the above example, this would result in the following global environment:

$$[\text{"odd"} \mapsto \mathsf{Vrecabs}\ css\ \text{"odd"}\ [], \text{"even"} \mapsto \mathsf{Vrecabs}\ css\ \text{"even"}\ []]$$
$$\text{where } css = [\text{"odd"} \mapsto [\langle 0\rangle \Rightarrow \langle \mathsf{False}\rangle, \langle \mathsf{Suc}\ n\rangle \Rightarrow \langle \mathsf{even}\ n\rangle],$$
$$\text{"even"} \mapsto [\langle 0\rangle \Rightarrow \langle \mathsf{True}\rangle, \langle \mathsf{Suc}\ n\rangle \Rightarrow \langle \mathsf{odd}\ n\rangle]]$$

Note that in the first line, the right-hand sides are values, but in *css*, they are expressions. The additional string argument of Vrecabs denotes the selected function. When evaluating an application of a recursive closure to an argument (β-reduction), the semantics adds all constituent functions of the closure to the environment used for recursive evaluation.

5 INTERMEDIATE SEMANTICS AND COMPILER PHASES

In this section, we will discuss the progression from de Bruijn based term language with its small-step semantics given in Figure 1a to the final CakeML semantics. The compiler starts out with terms of type term and applies multiple phases to eliminate features that are not present in the CakeML source language. Types term, nterm and pterm each have a small-step semantics only. Type sterm has a small-step and several intermediate big-step semantics that bridge the gap to CakeML. An overview of the intermediate semantics and compiler phases is depicted in Figure 3. The left-hand column gives an overview of the different phases. The right-hand column gives the types of the rule set and the semantics for each phase; you may want to skip it upon first reading.

5.1 Side conditions

All of the following semantics require some side conditions on the rule set. These conditions are purely syntactic. As an example we list the conditions for the correctness of the first compiler phase:

- Patterns must be linear, and constructors in patterns must be fully applied.
- Definitions must have at least one parameter on the left-hand side (§5.6).
- The right-hand side of an equation refers only to free variables occurring in patterns on the left-hand side and contain no dangling de Bruijn indices.
- There are no two defining equations $lhs = rhs_1$ and $lhs = rhs_2$ such that $rhs_1 \neq rhs_2$.

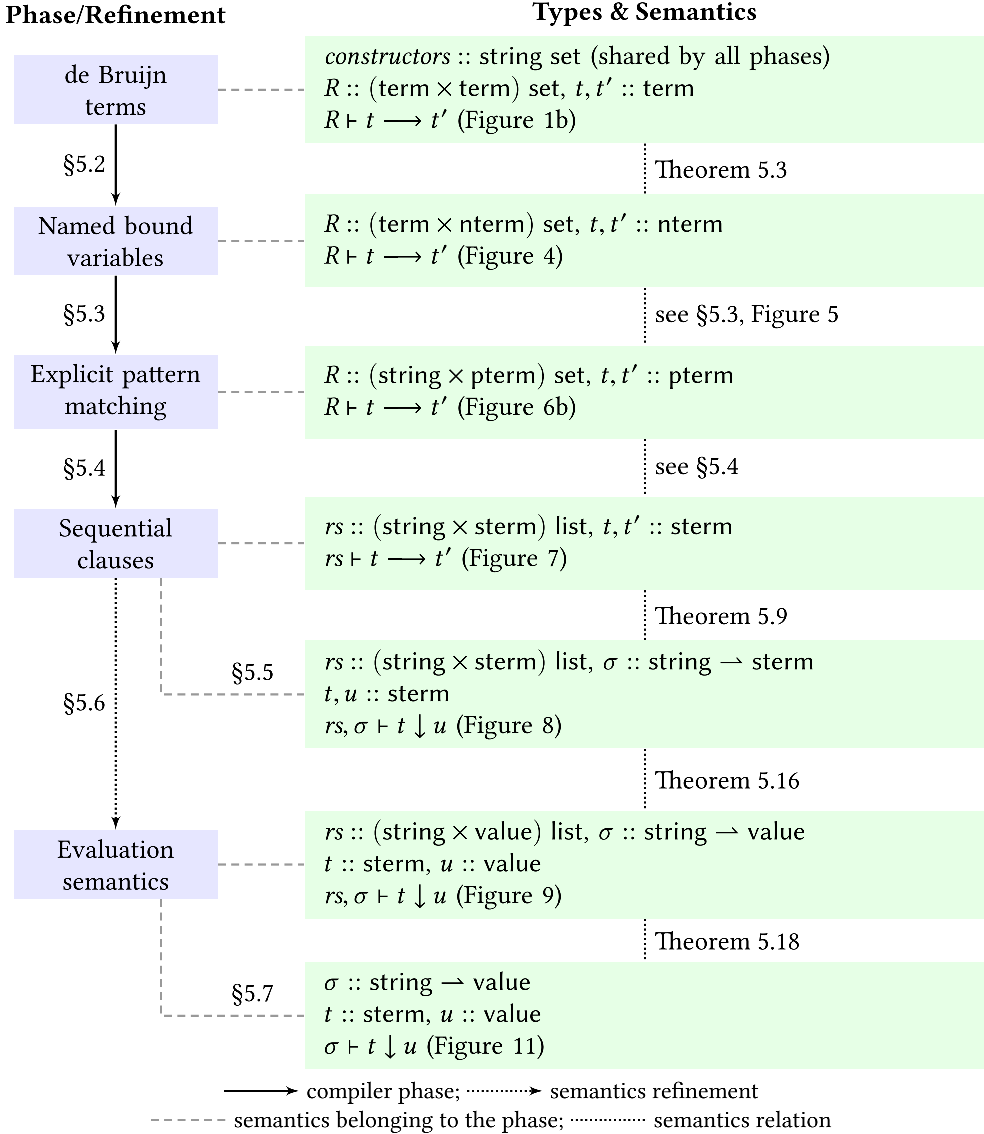

Fig. 3. Intermediate semantics and compiler phases

- For each pair of equations that define the same constant, their arity must be equal and their patterns must be compatible (§5.3).
- There is at least one equation.
- Variable names occurring in patterns must not overlap with constant names (§5.7).
- Any occurring constants must either be defined by an equation or be a constructor.

The conditions for the subsequent phases are sufficiently similar that we do not list them again.

$$\text{STEP}\ \frac{(lhs, rhs) \in R \qquad \text{match}\ lhs\ t = \text{Some}\ \sigma}{R \vdash t \longrightarrow \text{subst}\ \sigma\ rhs} \qquad \text{BETA}\ \frac{\text{closed}\ t'}{R \vdash (\Lambda x.\ t)\ \$\ t' \longrightarrow \text{subst}\ [x \mapsto t']\ \ t}$$

Fig. 4. Small-step semantics for type nterm with named bound variables

In the formalization, we use named contexts to fix the rules and assumptions on them (*locales* in Isabelle terminology [4]). Each phase has its own locale, together with a proof that after compilation, the preconditions of the next phase are satisfied. Correctness proofs assume the above conditions on R and similar conditions on the term that is reduced. For brevity, this is usually omitted in our presentation.

5.2 Naming bound variables: From term to nterm

Isabelle uses de Bruijn indices in the term language for the following two reasons: For substitution, there is no need to rename bound variables. Additionally, α-equivalent terms are equal. In implementations of programming languages, these advantages are not required: Typically, substitutions do not happen inside abstractions, and there is no notion of equality of functions. Therefore CakeML uses named variables and in this compilation step, we get rid of de Bruijn indices.

The "named" semantics is based on the nterm type (Figure 2). The rules that are changed from the original semantics (Figure 1b) are given in Figure 4 (Fun and Arg remain unchanged). Notably, β-reduction reuses the substitution function.

For the correctness proof, we need to establish a correspondence between terms and nterms. Translation from nterm to term is trivial: Replace bound variables by the number of abstractions between occurrence and where they were bound in, and keep free variables as they are. This function is called nterm_to_term.

The other direction is not unique and requires introduction of *fresh* names for bound variables. In our formalization, we have chosen to use a *monad* to produce these names. This function is called term_to_nterm.

Lemma 5.1 (Correctness of translation). *Let t be a* term *without dangling de Bruijn indices. Then* nterm_to_term (term_to_nterm t) $= t$.

The following sections will elaborate on the implementation and proof idea.

5.2.1 The fresh *monad.* Generation of fresh names in general can be thought of as picking a string that is not an element of a (finite) set of already existing names. For Isabelle, the *Nominal* framework [40, 41] provides support for reasoning over fresh names, but unfortunately, its definitions are not executable.

Instead, we chose to model generation of fresh names as a monad α fresh with the following primitive operations in addition to the monad operations:

$$\text{run} :: \alpha\ \text{fresh} \Rightarrow \text{string set} \Rightarrow \alpha$$
$$\text{fresh_name} :: \text{string fresh}$$

In our implementation, we have chosen to represent α fresh as roughly isomorphic to the state monad.

To avoid collisions of names further in the compilation pipeline, we additionally add the names of all constants to the set of known names.

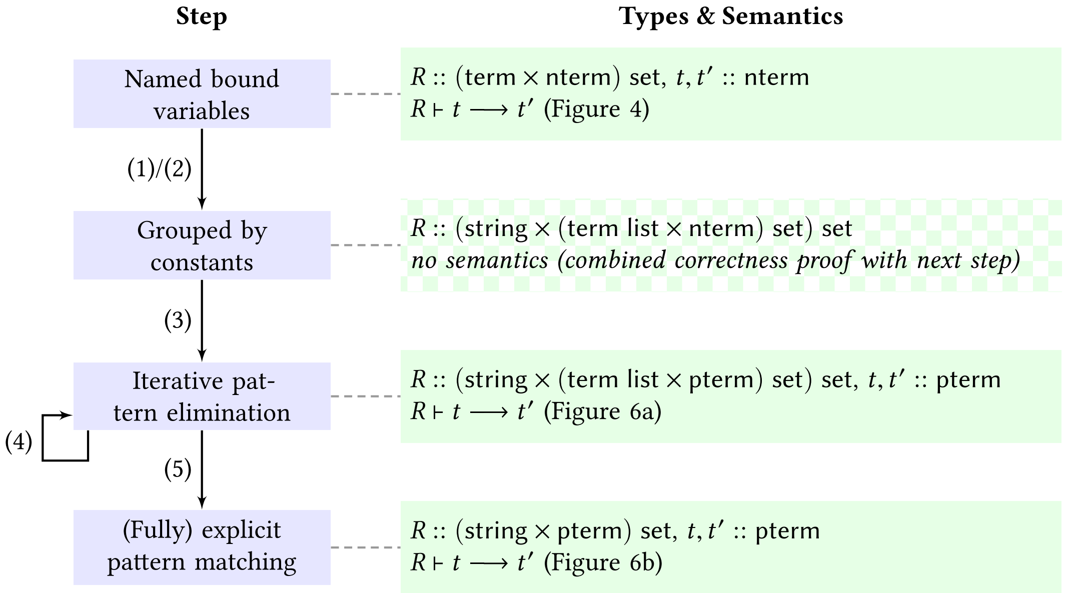

Fig. 5. Intermediate semantics, zoomed in on pattern elimination (step numbers refer to §5.3.3)

5.2.2 Implementation. The term_to_nterm function is implemented by structural recursion over the term, with an additional parameter Γ :: string list that records the context of fresh names generated at the enclosing abstractions, starting with $[]$. In the case of abstraction Λt, we invent a fresh name and add it to Γ for translating t.

Compilation of a rule set proceeds by translation of the right-hand side of all rules:

$$\text{compile } R = \{(p, \text{term_to_nterm } t) \mid (p, t) \in R\}$$

The left-hand side is left unchanged for two reasons: function match expects an argument of type term (see §4), and patterns do not contain abstractions or bound variables.

5.2.3 Correctness. We first prove a lemma about subst and β-reduction:

Lemma 5.2. $\text{nterm_to_term } \Gamma \; (\text{subst } [x \mapsto t'] \; t) = (\text{nterm_to_term } (x\#\Gamma) \; t)[\text{nterm_to_term } \Gamma \; t']$

Recall that $t[t']$ means β-reduction for the term type, i.e. with de Bruijn indices. Informally speaking, this means that we can either substitute t' for x in t and then translate the resulting nterm back to term, or translate back both t and t' and then perform β-reduction.

Theorem 5.3 (Correctness of compilation). *Assuming a step can be taken with the compiled rule set, it can be reproduced with the original rule set.*

$$\frac{\text{compile } R \vdash t \longrightarrow u \qquad \text{closed } t}{R \vdash \text{nterm_to_term } t \longrightarrow \text{nterm_to_term } u}$$

We prove this by induction over the semantics (Figure 4). The interesting case is Beta, which is immediately proved by the above lemma.

$$\textsc{Step}\ \frac{(f, R_f) \in R \qquad ([p_1, \ldots, p_n], rhs) \in R_f \qquad \text{match } (\text{Pconst } f \mathbin{\$} p_1 \mathbin{\$} \ldots \mathbin{\$} p_n)\ t = \text{Some } \sigma}{R \vdash t \longrightarrow \text{subst } \sigma\ rhs}$$

$$\textsc{Beta}\ \frac{(pat, rhs) \in C \qquad \text{match } pat\ t = \text{Some } \sigma \qquad \text{closed } t}{R \vdash (\Lambda\ C) \mathbin{\$} t \longrightarrow \text{subst } \sigma\ rhs}$$

(a) Combined implicit and explicit pattern matching

$$\textsc{Step'}\ \frac{(name, rhs) \in R}{R \vdash \text{Pconst } name \longrightarrow rhs}$$

(b) Modified Step rule for explicit-only pattern matching

Fig. 6. Small-step semantics for pterm with pattern matching

5.3 Explicit pattern matching: From nterm to pterm

Usually, functions in HOL are defined using *implicit* pattern matching, that is, the left-hand side of an equation is of the form $\langle \text{f}\ p_1\ \ldots\ p_n \rangle$. For any given function f, there may be multiple such equations. In this compilation step, we successively eliminate patterns from right-to-left until there is only one equation for f of the form $\langle \text{f} \rangle = \Lambda\ C$, where C is a set of clauses. After this step, the right-hand sides contain n nested abstractions, where n is the arity of f, i.e. the number of parameters of its defining equations.

Recall our running example (map). It has arity 2. We omit the brackets ⟨⟩ for brevity. First, the list parameter gets eliminated:

$$\begin{aligned}\text{map } f = \lambda\ [] &\Rightarrow [] \\ \mid x \mathbin{\#} xs &\Rightarrow f\ x \mathbin{\#} \text{map } f\ xs\end{aligned}$$

Finally, the function parameter gets eliminated:

$$\begin{aligned}\text{map} = \lambda\ f \Rightarrow (\lambda\ [] &\Rightarrow [] \\ \mid x \mathbin{\#} xs &\Rightarrow f\ x \mathbin{\#} \text{map } f\ xs)\end{aligned}$$

This has now arity 0 and is defined by a twice-nested abstraction.

5.3.1 Semantics. The intermediate semantics is given in Figure 6a (the Fun and Arg rules from previous semantics remain unchanged). We start out with a rule set R that allows both implicit and explicit pattern matching. Hence, both the Step and Beta functions need to call the match function. After elimination, only explicit pattern matching remains, which is reflected in Figure 6b. The modified Step rule merely replaces a constant by its definition, without taking arguments into account.

5.3.2 Elimination procedure. The elimination procedure can be described as an iterative matrix transformation. In a similar fashion as Slind's pattern matching compiler [37, §3.3.1], we view the set of defining equations of f as a matrix where each row is an equation.

$$\begin{array}{llcl} \text{f} & p_{1,1}\ \ldots\ p_{1,n} & = & rhs_1 \\ \text{f} & p_{2,1}\ \ldots\ p_{2,n} & = & rhs_2 \\ & & \vdots & \\ \text{f} & p_{m,1}\ \ldots\ p_{m,n} & = & rhs_m \end{array} \quad \rightsquigarrow \quad \left(\begin{array}{cccc|c} p_{1,1} & p_{1,2} & \cdots & p_{1,n} & rhs_1 \\ p_{2,1} & p_{2,2} & \cdots & p_{2,n} & rhs_2 \\ \vdots & \vdots & \ddots & \vdots & \vdots \\ p_{m,1} & p_{m,2} & \cdots & p_{m,n} & rhs_m \end{array}\right)$$

In each step, we group these equations by the initial $n-1$ patterns. More formally, we treat a row in the matrix as an $(n+1)$-tuple $(p_{i,1}, \ldots, p_{i,n}, rhs_i)$ and define an equivalence relation $\equiv_\text{p}$ such that

$$(p_{i,1}, \ldots, p_{i,n}, rhs_i) \equiv_\text{p} (p_{j,1}, \ldots, p_{j,n}, rhs_j) \Leftrightarrow (p_{i,1}, \ldots, p_{i,n-1}) = (p_{j,1}, \ldots, p_{j,n-1})$$

The new matrix is constructed from the set of equivalence classes. It consists of n columns and one row per equivalence class. The first $n-1$ columns contain the unique $p_{i,1}, \ldots, p_{i,n-1}$, whereas the last column is an abstraction with the set S_i of all $(p_{k,n}, rhs_k)$ in the equivalence class as the set of clauses.

$$\left(\begin{array}{cccc|c} p'_{1,1} & p'_{1,2} & \cdots & p'_{1,n-1} & \Lambda\ S_1 \\ p'_{2,1} & p'_{2,2} & \cdots & p'_{2,n-1} & \Lambda\ S_2 \\ \vdots & \vdots & \ddots & \vdots & \vdots \\ p'_{m',1} & p'_{m',2} & \cdots & p'_{m',n-1} & \Lambda\ S_{m'} \end{array}\right)$$

For the transformation to work, we need a strong assumption about the structure of the patterns p_i to avoid the following situation:

$$\begin{array}{llll} \mathsf{map} & f & [] & = [] \\ \mathsf{map} & g & (x \mathbin{\#} xs) & = g\ x \mathbin{\#} \mathsf{map}\ g\ xs \end{array}$$

Through elimination, this would turn into:

$$\begin{aligned} \mathsf{map} = \lambda\ f &\Rightarrow \big(\lambda\ [] \Rightarrow []\big) \\ \mid g &\Rightarrow \big(\lambda\ x \mathbin{\#} xs \Rightarrow f\ x \mathbin{\#} \mathsf{map}\ f\ xs\big) \end{aligned}$$

Even though the original equations were non-overlapping, we suddenly obtained an abstraction with two overlapping patterns. Slind observed a similar problem [37, §3.3.2] in his algorithm. Therefore, he only permits *uniform* equations, as defined by Wadler [34, §5.5]. Wadler also carries out the proof that the definition order of uniform equations does not matter. However, neither give a syntactic criterion of uniformity that does not require running their respective pattern compilation algorithms. The classic example for a non-uniform set of equations is the diagonal function:

fun diagonal :: bool ⇒ bool ⇒ bool ⇒ nat **where**
diagonal _ True False = 1
diagonal False _ True = 2
diagonal True False _ = 3

While our algorithm also cannot deal with this definition, we are able to formally characterize our requirements as a computable function on pairs of patterns:

fun pat_compatible :: term ⇒ term ⇒ bool **where**
pat_compatible (t_1 \$ t_2) (u_1 \$ u_2) ↔ pat_compatible t_1 u_1 ∧ (t_1 = u_1 ⟶ pat_compatible t_2 u_2)
pat_compatible t u ↔ (overlapping t u ⟶ t = u)

This compatibility constraint ensures that any two overlapping patterns (of the same column) $p_{i,k}$ and $p_{j,k}$ are equal and are thus appropriately grouped together in the elimination procedure. We require all defining equations of a constant to be mutually compatible. Equations violating this constraint will be flagged during embedding (§3), whereas the pattern elimination algorithm always succeeds.

Lemma 5.4 (Invariant). *Given a set of non-overlapping and pattern-compatible equations, the elimination procedure produces another set of non-overlapping and pattern-compatible equations.*

5.3.3 Implementation. After transformation to nterm, equations have the type term × nterm. Below we describe the stepwise transformation into a set of equations of type string × pterm. A diagram of the following algorithm is given in Figure 5.

(1) The left-hand side of each equation (of type term) is destructured into a tuple $(name, pats)$ :: string × term list, where *name* denotes the function symbol of the equation and *pats* the list of patterns.
(2) Equations with the same function symbol are grouped together. Different groups are processed separately. A single group has the type (term list × nterm) set. This corresponds to the initial matrix representation as depicted above.
(3) The right-hand side (of type nterm) can be trivially embedded into pterm: An nterm-abstraction $\Lambda x.\ t$ is translated to the pterm-abstraction $\Lambda\{\langle x\rangle \Rightarrow t\}$, i.e. a case abstraction with the single clause (Pvar x, t).
(4) For each function (i.e. for each group) that has an arity greater than zero, we apply the elimination procedure (compile_single :: (term list×pterm) set ⇒ (term list×pterm) set). This can be iterated until all functions have arity zero.
(5) A function of arity zero can be represented by a single pterm. The type of the rule set can thus be simplified to (string × pterm) set.

For correctness purposes, all but step (4) are trivial. For step (4), we first need to figure out what the correct correspondence relation is.

5.3.4 Correspondence relation. The statement of the semantic correctness property is more difficult than in the previous phase. The obvious property does not hold:

$$\frac{\text{compile_single } R \vdash t \longrightarrow u \qquad \text{closed } t}{R \vdash t \longrightarrow u}$$

Consider the map function again. After eliminating the right-most patterns, the defining equation of map is of the form $\langle \text{map } f = \lambda \ldots \rangle$, which means that we can rewrite the term $\langle \text{map id} \rangle$. However, we cannot reproduce this rewrite step in the original rule set, because the second argument is missing. Thus, we need to introduce a dedicated correspondence relation $\approx_\text{p}$ that takes extensionality and R into account. Therefore $\approx_\text{p}$ is implicitly parametrized by R. Note that extensionality is only needed if the translated term contains a Λ-abstraction.

Definition 5.5 (Left-deferred correspondence). Given a rule set R :: (string × pterm) set, the correspondence relation $\approx_\text{p}$ is defined inductively by the following rules:

$$\textsc{Const}\ \frac{}{\text{Pconst } n \approx_\text{p} \text{Pconst } n} \qquad \textsc{Var}\ \frac{}{\text{Pvar } n \approx_\text{p} \text{Pvar } n} \qquad \textsc{Comb}\ \frac{t_1 \approx_\text{p} t_2 \qquad u_1 \approx_\text{p} u_2}{t_1 \$ u_1 \approx_\text{p} t_2 \$ u_2}$$

$$\textsc{Abs}\ \frac{\text{rel}_\text{set}\ (\lambda(p_1, t_1)\ (p_2, t_2).\ p_1 = p_2 \wedge t_1 \approx_\text{p} t_2)\ C_1\ C_2}{\Lambda\ C_1 \approx_\text{p} \Lambda\ C_2}$$

$$\textsc{Defer}\ \frac{\begin{array}{c}\forall i.\ \text{closed } t_i \qquad (f, R_f) \in R \\ \text{arity } R_f > 0 \qquad \text{rel}_\text{set}\ (\lambda(p_1, t_1)\ (p_2, t_2).\ p_1 = p_2 \wedge t_1 \approx_\text{p} t_2)\ (\text{deferred } [t_1, \ldots, t_n]\ R_f)\ C\end{array}}{\text{Pconst } f \$ t_1 \$ \ldots \$ t_n \approx_\text{p} \Lambda\ C}$$

$$\begin{aligned}\text{deferred } ts\ R_f = \{&(p_{n+1}, \text{subst } \sigma\ rhs) \mid \\ &([p_1, p_2, \ldots, p_{n+1}], rhs) \in R_f \wedge \text{matchs } [p_1, \ldots, p_n]\ ts = \text{Some } \sigma\}\end{aligned}$$

We can illustrate the meaning of this relation based on the map example. Consider the matrix representation of the function:

$$R_{\mathsf{map}} = \left(\begin{array}{cc|c} \langle f \rangle & \langle [] \rangle & \langle [] \rangle \\ \langle f \rangle & \langle x \mathbin{\#} xs \rangle & \langle f\ x \mathbin{\#} \mathsf{map}\ f\ xs \rangle \end{array} \right)$$

A single compilation step turns that into:

$$R'_{\mathsf{map}} = \left(\begin{array}{c|c} \langle f \rangle & \left\langle \begin{array}{l} \lambda\ [] \Rightarrow [] \\ \mid x \mathbin{\#} xs \Rightarrow f\ x \mathbin{\#} \mathsf{map}\ f\ xs \end{array} \right\rangle \end{array} \right)$$

In this modified rule set, the term $\langle \mathsf{map}\ (\lambda y.\ y) \rangle$ can be rewritten. In the original rule set, it cannot. With the Defer rule, we can still relate the un-reduced term to the reduced term:

$$\langle \mathsf{map}\ (\lambda y.\ y) \rangle \approx_{\mathrm{p}} \left\langle \begin{array}{l} \lambda\ [] \Rightarrow [] \\ \mid x \mathbin{\#} xs \Rightarrow (\lambda y.\ y)\ x \mathbin{\#} \mathsf{map}\ (\lambda y.\ y)\ xs \end{array} \right\rangle$$

The Defer rule can be explained by examining the deferred function. Given a function application for n parameters $\langle \mathsf{f}\ t_1\ \ldots\ t_n \rangle$ of a function with arity $n+1$, it selects all defining equations that match these arguments. Each of these equations $(\langle \mathsf{f}\ p_1\ \ldots\ p_n\ p_{n+1} \rangle, \langle t \rangle)$ carries an additional pattern p_{n+1} for the $(n+1)$st argument which has to be supplied eventually. From these equations, we construct a Λ-abstraction comprising pairs $(\langle p_{n+1} \rangle, \mathsf{subst}\ \sigma\ \langle t \rangle)$, where σ is the result of matching the initial n patterns.

In other words, the left-deferred correspondence can also be described as a *right-extensional* correspondence. We exploit the idea that we can relate f and g if for all x, $f\ x$ is related to $g\ x$. It is "right-extensional" in the sense that g needs to be a Λ-abstraction.

5.3.5 *Correctness.*

Theorem 5.6 (Correctness).

$$\frac{\mathsf{compile_single}\ R \vdash u \longrightarrow u' \qquad t \approx_{\mathrm{p}} u \qquad \mathsf{closed}\ t}{\exists t'.\ R \vdash t \longrightarrow^{*} t' \wedge t' \approx_{\mathrm{p}} u'}$$

Proof. The main idea is that the rewrite steps for a term t that are possible in compile_single R are captured in the deferred set. We prove this by rule induction on the semantics (Figure 6a). The interesting cases are Step and Beta.

Step Consider the term $\langle \mathsf{f}\ t_1\ \ldots\ t_n \rangle$. As explained above, the new rule set may be able to perform a rewrite step which cannot be reproduced in the original set. Whether or not that is the case depends on the arity of the constant in the original rule set. If f already had arity zero, nothing changed during compilation. If not, we must use the Defer rule to relate the term $\langle \mathsf{f}\ t_1\ \ldots\ t_n \rangle$ to a case abstraction. In essence, this case is the correctness proof of the deferred set.

Beta Here, we need to reproduce a β-reduction of a term of the form $\Lambda\ C\ \$\ u$. We also know that there is a term t such that $t \approx_{\mathrm{p}} \Lambda\ C$. We can apply rule inversion on this fact. The proof proceeds by case analysis of the two possible rules Ext and Defer. Ext is trivial, because it reveals that t has an identical structure to $\Lambda\ C$. Defer requires the inverse direction of the Step case, i.e. the completeness of the deferred set.

Note that in the conclusion of the theorem the reflexive-transitive closure of $\longrightarrow$ is used. But in both Step and Beta, we would only need reflexive closure, because $\longrightarrow$ is applied zero times or once. □

$$\textsc{Step}\ \frac{(\mathit{name}, \mathit{rhs}) \in R}{R \vdash \text{Sconst}\ \mathit{name} \longrightarrow \mathit{rhs}} \qquad \textsc{Beta}\ \frac{\text{first_match}\ \mathit{cs}\ t = \text{Some}\ (\sigma, \mathit{rhs}) \qquad \text{closed}\ t}{R \vdash (\Lambda\ \mathit{cs})\ \$\ t \longrightarrow \text{subst}\ \sigma\ \mathit{rhs}}$$

Fig. 7. Small-step semantics for type sterm

$$\textsc{Const}\ \frac{(\mathit{name}, \mathit{rhs}) \in \mathit{rs}}{\mathit{rs}, \sigma \vdash \text{Sconst}\ \mathit{name} \downarrow \mathit{rhs}} \qquad \textsc{Var}\ \frac{\sigma\ \mathit{name} = \text{Some}\ v}{\mathit{rs}, \sigma \vdash \text{Svar}\ \mathit{name} \downarrow v}$$

$$\textsc{Abs}\ \frac{}{\mathit{rs}, \sigma \vdash \Lambda\ \mathit{cs} \downarrow \Lambda\,[(\mathit{pat}, \text{subst}\ (\sigma - \text{frees}\ \mathit{pat})\ t \mid (\mathit{pat}, t) \leftarrow \mathit{cs}]}$$

$$\textsc{Comb}\ \frac{\begin{array}{c}\mathit{rs}, \sigma \vdash t \downarrow \Lambda\ \mathit{cs} \\ \mathit{rs}, \sigma \vdash u \downarrow u' \qquad \text{first_match}\ \mathit{cs}\ u' = \text{Some}\ (\sigma', \mathit{rhs}) \qquad \mathit{rs}, \sigma +\!\!+\ \sigma' \vdash \mathit{rhs} \downarrow v\end{array}}{\mathit{rs}, \sigma \vdash t\ \$\ u \downarrow v}$$

$$\textsc{Constr}\ \frac{\mathit{name} \in \mathit{constructors} \qquad \mathit{rs}, \sigma \vdash t_1 \downarrow u_1 \qquad \cdots \qquad \mathit{rs}, \sigma \vdash t_n \downarrow u_n}{\mathit{rs}, \sigma \vdash \text{Sconst}\ \mathit{name}\ \$\ t_1\ \$ \ldots \$\ t_n \downarrow \text{Sconst}\ \mathit{name}\ \$\ u_1\ \$ \ldots \$\ u_n}$$

Fig. 8. Big-step semantics for type sterm

5.4 Sequentialization: From pterm to sterm

The semantics of pterm and sterm differ only in rule Step and Beta. Figure 7 shows the modified rules, where instead of any matching clause the first matching clause in a case abstraction is picked. For the correctness proof, the order of clauses does not matter: we only need to prove that a step taken in the sequential semantics can be reproduced in the unordered semantics. As long as no rules are dropped, this is trivially true. For that reason, the compiler orders the clauses lexicographically. At the same time the rules are also converted from type (string $\times$ pterm) set to (string $\times$ sterm) list. Below, *rs* will always denote a list of the latter type.

Note that this semantics only sequentializes pattern matching: rewriting may still be non-deterministic because terms may contain multiple redexes.

5.5 Big-step semantics for sterm

This big-step semantics for sterm is not a compiler phase but moves towards the desired evaluation semantics. In this first step, we reuse the sterm type for evaluation results, instead of evaluating to the separate type value. This allows us to ignore environment capture in closures for now.

All previous $\longrightarrow$ relations were parametrized by a rule set. Now the big-step predicate is of the form $\mathit{rs}, \sigma \vdash t \downarrow t'$ where σ :: string $\rightharpoonup$ sterm is a variable environment.

This semantics also introduces the distinction between *constructors* and *defined constants.* If C is a constructor, the term $\langle \text{C}\ t_1\ \ldots\ t_n \rangle$ is evaluated to $\langle \text{C}\ t'_1\ \ldots\ t'_n \rangle$ where the t'_i are the results of evaluating the t_i.

The full set of rules is shown in Figure 8. They deserve a short explanation:

Const Constants are retrieved from the rule set *rs*.

Var Variables are retrieved from the environment σ.

Abs In order to achieve the intended invariant, abstractions are evaluated to their fully substituted form.

Comb Function application $t\ \$\ u$ first requires evaluation of t into an abstraction $\Lambda\ \mathit{cs}$ and evaluation of u into an arbitrary term u'. Afterwards, we look for a clause matching u' in *cs*,

$$\textsc{Const}\ \frac{(name, rhs) \in rs}{rs, \sigma \vdash \text{Sconst}\ name \downarrow rhs} \qquad \textsc{Var}\ \frac{\sigma\ name = \text{Some}\ v}{rs, \sigma \vdash \text{Svar}\ name \downarrow v} \qquad \textsc{Abs}\ \frac{}{rs, \sigma \vdash \Lambda\ cs \downarrow \text{Vabs}\ cs\ \sigma}$$

$$\textsc{Comb}\ \frac{\begin{array}{c} rs, \sigma \vdash t \downarrow \text{Vabs}\ cs\ \sigma' \\ rs, \sigma \vdash u \downarrow v \qquad \text{first_match}\ cs\ v = \text{Some}\ (\sigma'', rhs) \qquad rs, \sigma' \mathbin{+\!\!+} \sigma'' \vdash rhs \downarrow v' \end{array}}{rs, \sigma \vdash t\ \$\ u \downarrow v'}$$

$$\textsc{RecComb}\ \frac{\begin{array}{c} rs, \sigma \vdash t \downarrow \text{Vrecabs}\ css\ name\ \sigma' \qquad css\ name = \text{Some}\ cs \\ rs, \sigma \vdash u \downarrow v \qquad \text{first_match}\ cs\ v = \text{Some}\ (\sigma'', rhs) \qquad rs, \sigma' \mathbin{+\!\!+} \sigma'' \vdash rhs \downarrow v' \end{array}}{rs, \sigma \vdash t\ \$\ u \downarrow v'}$$

$$\textsc{Constr}\ \frac{name \in constructors \qquad rs, \sigma \vdash t_1 \downarrow v_1 \qquad \cdots \qquad rs, \sigma \vdash t_n \downarrow v_n}{rs, \sigma \vdash \text{Sconst}\ name\ \$\ t_1\ \$ \ldots \$\ t_n \downarrow \text{Vconstr}\ name\ [v_1, \ldots, v_n]}$$

Fig. 9. Evaluation semantics from sterm to value

which produces a local variable environment σ', possibly overwriting existing variables in σ. Finally, we evaluate the right-hand side of the clause with the combined global and local variable environment.

Constr For a constructor application $\langle \text{C}\ t_1\ \ldots \rangle$, evaluate all t_i. The set *constructors* is an implicit parameter of the semantics.

Lemma 5.7 (Closedness invariant). *If σ contains only closed terms,* frees $t \subseteq$ dom σ *and* $rs, \sigma \vdash t \downarrow t'$*, then t' is closed.*

Correctness of the big-step w.r.t. the small-step semantics is proved easily by induction on the former:

Lemma 5.8. *For any closed environment σ satisfying* frees $t \subseteq$ dom σ*,*

$$rs, \sigma \vdash t \downarrow u \rightarrow rs \vdash \text{subst}\ \sigma\ t \longrightarrow^* u$$

By setting $\sigma = []$, we obtain:

Theorem 5.9 (Correctness). $rs, [] \vdash t \downarrow u \wedge \text{closed}\ t \rightarrow rs \vdash t \longrightarrow^* u$

5.6 Evaluation semantics: Refining sterm to value

At this point, we introduce the concept of values into the semantics, while still keeping the rule set (for constants) and the environment (for variables) separate. The evaluation rules are specified in Figure 9 and represent a departure from the original rewriting semantics: a term does not evaluate to another term but to an object of a different type, a value. We still use $\downarrow$ as notation, because big-step and evaluation semantics can be disambiguated by their types.

The evaluation model itself is fairly straightforward. As explained in §4.5, abstraction terms are evaluated to closures capturing the current variable environment. Note that at this point, recursive closures are not treated differently from non-recursive closures. In a later stage, when *rs* and σ are merged, this distinction becomes relevant.

We will now explain each rule that has changed from the previous semantics:

Abs Abstraction terms are evaluated to a closure capturing the current environment.

datatype pat = Patvar string | Patconstr string (pat list)

mk_pat (Const *name* $ p_1 $ ... $ p_n) = Patconstr *name* (map mk_pat [$p_1, \ldots, p_n$])
mk_pat (Free *name*) = Patvar *name*

Fig. 10. Proper patterns

Comb As before, in an application $t \ \$ \ u$, t must evaluate to a closure Vabs $cs\ \sigma'$. The evaluation result of u is then matched against the clauses *cs*, producing an environment σ''. The right-hand side of the clause is then evaluated using $\sigma' +\!\!+ \sigma''$; the original environment σ is effectively discarded.

RecComb Similar as above. Finding the matching clause is a two-step process: First, the appropriate clause list is selected by name of the currently active function. Then, matching is performed.

Constr As before, for an n-ary application $\langle \mathrm{C}\ t_1 \ \ldots \rangle$, where C is a data constructor, we evaluate all t_i. The result is a Vconstr value.

5.6.1 Conversion between sterm *and* value*.* To establish a correspondence between evaluating a term to an sterm and to a value, we apply the same trick as in §5.2. Instead of specifying a complicated relation, we translate value back to sterm: simply apply the substitutions in the captured environments to the clauses.

fun list_comb :: sterm ⇒ sterm list ⇒ sterm **where**
list_comb *f* [] = *f*
list_comb *f* (*t* # *ts*) = list_comb (*f* $ *t*) *ts*

fun value_to_sterm :: value ⇒ sterm **where**
value_to_sterm (Vconstr *name vs*) = list_comb (Sconst *name*) (map value_to_sterm *vs*)
value_to_sterm (Vabs *cs* σ) = Λ [(*pat*, subst (map value_to_sterm σ − frees pat) *t*)) | (*pat*, *t*) ← *cs*]
value_to_sterm (Vrecabs *css name* σ) = value_to_sterm (Vabs (*css name*) σ)

The translation rules for Vabs and Vrecabs are intentionally similar to the Abs rule from the big-step semantics (Figure 8). Roughly speaking, the big-step semantics always keeps terms fully substituted, whereas the evaluation semantics defers substitution.

Similarly to Lemma 5.1, we can also define a function sterm_to_value :: sterm ⇒ value and prove that one function is the inverse of the other. Before we can state that property, it is important to observe that it does not hold for arbitrary terms, but just for *value terms*. These are characterized by the codomain of the value_to_sterm function. But we can also prove that the big-step semantics only produces value terms.

Lemma 5.10 (Correctness of translation). value_to_sterm (sterm_to_value t) $= t$ *for all value terms* t.

5.6.2 Matching. The value type, instead of using binary function application as all other term types, uses n-ary constructor application. This introduces a conceptual mismatch between (binary) patterns and values. To make the proofs easier, we introduce an intermediate type of n-ary patterns (Figure 10). The function mk_pat :: term ⇒ pat converts from binary to n-ary patterns. For readability, we show only its informal specification.

The Comb and RecComb rules use a variant of the first_match function that converts the term to pat first (with mk_pat) and uses vmatch :: pat ⇒ value ⇒ (string ⇀ value) option for matching. Note that this merely simplifies proofs: the intermediate pattern type can be optimized away by fusing the vmatch and mk_pat functions.

Lemma 5.11 (*n*-ary vs. binary patterns). *For all linear patterns, the result of* vmatch *corresponds (w.r.t.* value_to_sterm*) to* match. *Formally:*

$$\text{linear } p \longrightarrow \text{rel}_{\text{option}}\ (\text{rel}_{\text{mapping}}\ (\lambda v\ t.\ t = \text{value_to_sterm } v))\ \sigma_1\ \sigma_2$$
$$\textit{where } \sigma_1 = \text{vmatch } (\text{mk_pat } p)\ v$$
$$\sigma_2 = \text{match } p\ (\text{value_to_sterm } v)$$

5.6.3 Correctness. The correctness proof requires a number of interesting lemmas.

Lemma 5.12 (Substitution before evaluation). *Assuming that a term t can be evaluated to a value u given a closed environment σ, it can be evaluated to the same value after substitution with a sub-environment σ'. Formally: $rs, \sigma \vdash t \downarrow u \wedge \sigma' \subseteq \sigma \longrightarrow rs, \sigma \vdash \text{subst } \sigma'\ t \downarrow u$*

This justifies the “pre-substitution” exhibited by the Abs rule in the big-step semantics in contrast to the environment-capturing Abs rule in the evaluation semantics.

Lemma 5.13 (Environment coincidence). *Let σ, σ' be two closed environments whose mappings coincide on the set S. Formally: $S \subseteq \text{dom } \sigma \wedge \forall a \in S.\ \sigma\ a = \sigma'\ a$. Then, if* frees $t \subseteq S$*, evaluation in both environments yields the same result: $rs, \sigma \vdash t \downarrow u \leftrightarrow rs, \sigma' \vdash t \downarrow u$.*

The lemma itself is obvious. Its proof is mainly technical; it requires set-theoretic reasoning about domains of mappings.

Theorem 5.14 (Correctness). *Let σ be a closed environment and t a term which only contains free variables in* dom σ. *Then, an evaluation to a value $rs, \sigma \vdash t \downarrow v$ can be reproduced in the big-step semantics as $rs', \text{map value_to_sterm } \sigma \vdash t \downarrow \text{value_to_sterm } v$, where $rs' = [(name, \text{value_to_sterm } rhs) \mid (name, rhs) \leftarrow rs]$.*

Proof. By induction over the evaluation semantics. The interesting cases are Comb and RecComb. Both are roughly identical, save for the additional complication that RecComb has an extra selection step to find the clause set. We will omit that for brevity and focus on the Comb case. In this case, we need to prove that function application behaves the same way in both semantics. The major difference between the two different Comb rules is that in the evaluation semantics, the evaluation of t to a closure cs reveals a hidden environment σ' that is subsequently used in the evaluation of rhs (where $(pat, rhs) \in cs$). That hidden environment may bear no relation to the active environment σ which is used in the big-step semantics. Hence, we use a trick: We “pre-substitute” in rhs (Lemma 5.12) and can then use Lemma 5.13 to replace σ' for σ (side-conditions omitted):

$$\dfrac{\dfrac{\dfrac{\dfrac{}{rs', \text{map value_to_sterm } (\sigma' +\!\!+ \sigma'') \vdash rhs \downarrow \text{value_to_sterm } v'}\ \text{IH}}{rs', \text{map value_to_sterm } (\sigma' +\!\!+ \sigma'') \vdash rhs_{\sigma'} \downarrow \text{value_to_sterm } v'}\ \text{Lemma 5.12}}{rs', \text{map value_to_sterm } (\sigma +\!\!+ \sigma'') \vdash rhs_{\sigma'} \downarrow \text{value_to_sterm } v'}\ \text{Lemma 5.13}}{rs', \text{map value_to_sterm } \sigma \vdash t\ \$\ u \downarrow \text{value_to_sterm } v'}\ \textsc{Comb}$$

where $rhs_{\sigma'} = \text{subst } (\text{map value_to_sterm } \sigma' - \text{frees } pat)\ rhs$. We specifically chose this value for $rhs_{\sigma'}$ to be able to combine both lemmas and to coincide with the definition of value_to_sterm. □

5.6.4 Instantiating the correctness theorem. The correctness theorem states that, for any given evaluation of a term t with a given environment rs, σ containing values, we can reproduce that evaluation in the big-step semantics using a derived list of rules rs' and an environment σ' containing sterms that are generated by the value_to_sterm function. But recall the diagram in Figure 3. In our scenario, we start with a given rule set of sterms (that has been compiled from a rule set of terms). Hence, the correctness theorem only deals with the opposite direction.

It remains to construct a suitable rs such that applying value_to_sterm to it yields the given sterm rule set. We can exploit the side condition (§5.1) that all bindings define functions, not constants:

Definition 5.15 (Global clause set). The mapping global_css :: string $\rightharpoonup$ ((term $\times$ sterm) list) is obtained by stripping the Sabs constructors from all definitions and converting the resulting list to a mapping.

For each definition with name f we define a corresponding term $v_f = \text{Vrecabs global_css } f\ []$. In other words, each function is now represented by a recursive closure bundling all functions. Applying value_to_sterm to v_f returns the original definition of f. Let rs denote the original sterm rule set and rs_v the environment mapping all f's to the v_f's.

The variable environments σ and σ' can safely be set to the empty mapping, because top-level terms are evaluated without any free variable bindings.

Corollary 5.16 (Correctness). $rs_\text{v}, [] \vdash t \downarrow v \longrightarrow rs, [] \vdash t \downarrow \text{value_to_sterm } v$

Note that this step was not part of the compiler (although rs_v is computable) but it is a refinement of the semantics to support a more modular correctness proof.

Example. Recall the odd and even example from §4.5. After compilation to sterm, the rule set looks like this:

$$\begin{aligned} rs = \{(\text{"odd"}, \text{Sabs } [\langle 0 \rangle \Rightarrow \langle \text{False} \rangle, \langle \text{Suc } n \rangle \Rightarrow \langle \text{even } n \rangle]), \\ (\text{"even"}, \text{Sabs } [\langle 0 \rangle \Rightarrow \langle \text{True} \rangle, \langle \text{Suc } n \rangle \Rightarrow \langle \text{odd } n \rangle])\} \end{aligned}$$

This can be easily transformed into the following global clause set:

$$\begin{aligned} \text{global_css} = [\text{"odd"} \mapsto [\langle 0 \rangle \Rightarrow \langle \text{False} \rangle, \langle \text{Suc } n \rangle \Rightarrow \langle \text{even } n \rangle], \\ \text{"even"} \mapsto [\langle 0 \rangle \Rightarrow \langle \text{True} \rangle, \langle \text{Suc } n \rangle \Rightarrow \langle \text{odd } n \rangle]] \end{aligned}$$

Finally, rs_v is computed by creating a recursive closure for each function:

$$rs_\text{v} = [\text{"odd"} \mapsto \text{Vrecabs global_css "odd" } [], \text{"even"} \mapsto \text{Vrecabs global_css "even" } []]$$

5.7 Evaluation with recursive closures

CakeML distinguishes between non-recursive and recursive closures [29]. This distinction is also present in the value type. In this step, we will conflate variables with constants which necessitates a special treatment of recursive closures. Therefore we introduce a new predicate $\sigma \vdash t \downarrow v$ in Figure 11 (in contrast to the previous $rs, \sigma \vdash t \downarrow v$). We examine the rules one by one:

Const/Var Constant definition and variable values are both retrieved from the same environment σ. We have opted to keep the distinction between constants and variables in the sterm type to avoid the introduction of another term type.

Abs Identical to the previous evaluation semantics. Note that evaluation never creates recursive closures at run-time (only at compile-time, see §5.6.4). Anonymous functions, e.g. in the term $\langle \text{map } (\lambda x.\ x) \rangle$, are evaluated to non-recursive closures.

$$\textsc{Const}\ \frac{\mathit{name} \notin \mathit{constructors} \qquad \sigma\ \mathit{name} = \mathsf{Some}\ v}{\sigma \vdash \mathsf{Sconst}\ \mathit{name} \downarrow v}$$

$$\textsc{Var}\ \frac{\sigma\ \mathit{name} = \mathsf{Some}\ v}{\sigma \vdash \mathsf{Svar}\ \mathit{name} \downarrow v} \quad \textsc{Abs}\ \frac{}{\sigma \vdash \Lambda\ cs \downarrow \mathsf{Vabs}\ cs\ \sigma}$$

$$\textsc{Comb}\ \frac{\begin{array}{c}\sigma \vdash t \downarrow \mathsf{Vabs}\ cs\ \sigma' \\ \sigma \vdash u \downarrow v \qquad \mathsf{first_match}\ cs\ v = \mathsf{Some}\ (\sigma'', \mathit{rhs}) \qquad \sigma' \mathbin{+\!\!+} \sigma'' \vdash \mathit{rhs} \downarrow v'\end{array}}{\sigma \vdash t\ \$\ u \downarrow v'}$$

$$\textsc{RecComb}\ \frac{\begin{array}{c}\sigma \vdash t \downarrow \mathsf{Vrecabs}\ css\ \mathit{name}\ \sigma' \qquad css\ \mathit{name} = \mathsf{Some}\ cs \qquad \sigma \vdash u \downarrow v \\ \mathsf{first_match}\ cs\ v = \mathsf{Some}\ (\sigma'', \mathit{rhs}) \qquad \sigma' \mathbin{+\!\!+} \mathsf{mk_rec_env}\ css\ \sigma' \mathbin{+\!\!+} \sigma'' \vdash \mathit{rhs} \downarrow v'\end{array}}{\sigma \vdash t\ \$\ u \downarrow v'}$$

$$\textsc{Constr}\ \frac{\mathit{name} \in \mathit{constructors} \qquad \sigma \vdash t_1 \downarrow v_1 \qquad \cdots \qquad \sigma \vdash t_n \downarrow v_n}{\sigma \vdash \mathsf{Sconst}\ \mathit{name}\ \$\ t_1\ \$ \ldots \$\ t_n \downarrow \mathsf{Vconstr}\ \mathit{name}\ [v_1, \ldots, v_n]}$$

Fig. 11. ML-style evaluation semantics

COMB Identical to the previous evaluation semantics.

RECCOMB Almost identical to the evaluation semantics. For each function $(\mathit{name}, cs) \in css$, a new recursive closure Vrecabs $css\ \mathit{name}\ \sigma'$ is created and inserted into the environment. This ensures that after the first call to a recursive function, the function itself is present in the environment to be called recursively, without having to introduce coinductive environments.

CONSTR Identical to the evaluation semantics.

5.7.1 Conflating constants and variables. By merging the rule set *rs* with the variable environment σ, it becomes necessary to discuss possible clashes. Previously, the syntactic distinction between Svar and Sconst meant that $\langle x \rangle$ and $\langle \mathrm{x} \rangle$ are not ambiguous: all semantics up to the evaluation semantics clearly specify where to look for the substitute. This is not the case in functional languages where functions and variables are not distinguished syntactically.

Instead, we rely on the fact that the initial rule set only defines constants. All variables are introduced by matching before β-reduction (that is, in the COMB and RECCOMB rules). The ABS rule does not change the environment. Hence it suffices to assume that variables in patterns must not overlap with constant names (see §5.1).

5.7.2 Correspondence relation. Both constant definitions and values of variables are recorded in a single environment σ. This also applies to the environment contained in a closure. The correspondence relation thus needs to take a different sets of bindings in closures into account.

Hence, we define a relation $\approx_{\mathrm{v}}$ that is implicitly parametrized on the rule set *rs* and compares environments. We call it *right-conflating*, because in a correspondence $v \approx_{\mathrm{v}} u$, any bound environment in u is thought to contain both variables and constants, whereas in v, any bound environment contains only variables.

Definition 5.17 (Right-conflating correspondence). We define $\approx_{\mathrm{v}}$ coinductively as follows:

$$\frac{v_1 \approx_{\mathrm{v}} u_1 \qquad \cdots \qquad v_n \approx_{\mathrm{v}} u_n}{\mathsf{Vconstr}\ \mathit{name}\ [v_1, \ldots, v_n] \approx_{\mathrm{v}} \mathsf{Vconstr}\ \mathit{name}\ [u_1, \ldots, u_n]}$$

$$\frac{\forall x \in \mathsf{frees}\ cs.\ \sigma_1\ x \approx_{\mathrm{v}} \sigma_2\ x \qquad \forall x \in \mathsf{consts}\ cs.\ rs\ x \approx_{\mathrm{v}} \sigma_2\ x}{\mathsf{Vabs}\ cs\ \sigma_1 \approx_{\mathrm{v}} \mathsf{Vabs}\ cs\ \sigma_2}$$

$$\frac{\begin{array}{c}\forall cs \in \mathsf{range}\ css.\ \forall x \in \mathsf{frees}\ cs.\ \sigma_1\ x \approx_{\mathrm{v}} \sigma_2\ x \\ \forall cs \in \mathsf{range}\ css.\ \forall x \in \mathsf{consts}\ cs.\ \sigma_1\ x \approx_{\mathrm{v}} (\sigma_2 +\!\!+ \mathsf{mk_rec_env}\ css\ \sigma_2)\ x\end{array}}{\mathsf{Vrecabs}\ css\ name\ \sigma_1 \approx_{\mathrm{v}} \mathsf{Vrecabs}\ css\ name\ \sigma_2}$$

Consequently, $\approx_{\mathrm{v}}$ is not reflexive.

5.7.3 Correctness. The correctness lemma is straightforward to state:

THEOREM 5.18 (CORRECTNESS). *Let σ be an environment, t be a closed term and v a value such that $\sigma \vdash t \downarrow v$. If for all constants x occurring in t, $rs\ x \approx_{\mathrm{v}} \sigma\ x$ holds, then there is an u such that $rs, [] \vdash t \downarrow u$ and $u \approx_{\mathrm{v}} v$.*

As usual, the rather technical proof proceeds via induction over the semantics (Figure 11). It is important to note that the global clause set construction (§5.6.4) satisfies the preconditions of this theorem:

LEMMA 5.19. *If name is the name of a constant in rs, then* $\mathsf{Vrecabs}\ \mathsf{global_css}\ name\ [] \approx_{\mathrm{v}} \mathsf{Vrecabs}\ \mathsf{global_css}\ name\ []$.

Because $\approx_{\mathrm{v}}$ is defined coinductively, the proof of this precondition proceeds by coinduction.

5.8 CakeML

CakeML is a verified implementation of a subset of Standard ML [23, 38]. It comprises a parser, type checker, formal semantics and backend for machine code. The semantics has been formalized in Lem [28], which allows export to Isabelle theories.

Our compiler targets CakeML's abstract syntax tree. However, we do not make use of certain CakeML features; notably mutable cells, modules, and literals. We have derived a smaller, executable version of the original CakeML semantics, called *CupCakeML*, together with an equivalence proof. The correctness proof of the last compiler phase establishes a correspondence between CupCakeML and the final semantics of our compiler pipeline.

For the correctness proof of the CakeML compiler, its authors have extracted the Lem specification into HOL4 theories [1]. In our work, we directly target CakeML abstract syntax trees (thereby bypassing the parser) and use its big-step semantics, which we have extracted into Isabelle.[3]

5.8.1 CupCakeML. The core of CakeML's big-step semantics is organized in three mutually recursive inductive predicates (evaluate, evaluate_list, evaluate_match) with a total of over thirty rules. It is that large because it has to deal with exceptions, modules, step counters, and other features that their compiler supports, but we do not need. For that reason, we have defined the subset *CupCakeML* which only allows the syntactic forms known from the sterm type.

We define a smaller big-step semantics with a single inductive predicate and twelve remaining rules. We then prove equivalence to the original semantics, that is, both correctness and completeness, under the assumption that the expression and initial environments are in the supported fragment. We also prove that the supported fragment is closed under the relevant functions and relations.

[3]based on a repository snapshot from March 27, 2017 (0c48672)

5.8.2 Conversion from sterm *to* exp. After the series of translations described in the earlier sections, our terms are syntactically close to CakeML's terms (Cake.exp). The only remaining differences are outlined below:

- CakeML does not combine abstraction and pattern matching. For that reason, we have to translate $\Lambda\ [p_1 \Rightarrow t_1, \ldots]$ into $\Lambda x.$ **case** x **of** $p_1 \Rightarrow t_1 \mid \ldots$, where x is a fresh variable name. We reuse the fresh monad to obtain a bound variable name. Note that it is not necessary to thread through already created variable names, only existing names. The reason is simple: a generated variable is bound and then immediately used in the body. Shadowing it somewhere in the body is not problematic.
- CakeML has two distinct syntactic categories for identifiers (that can represent variables or functions) and data constructors. Our term types however have two distinct syntactic categories for constants (that can represent functions or data constructors) and variables. The necessary prerequisites to deal with this are already present in the ML-style evaluation semantics (§5.7) which conflates constants and variables, but has a dedicated Constr rule for data constructors.

5.8.3 Types. During embedding (§3), all type information is erased. Yet, CakeML performs some limited form of type checking at run-time: constructing and matching data must always be fully applied. That is, data constructors must always occur with all arguments supplied on right-hand and left-hand sides.

Fully applied constructors in terms can be easily guaranteed by simple pre-processing. For patterns however, this must be ensured throughout the compilation pipeline; it is (like other syntactic constraints) another side condition imposed on the rule set (§5.1).

The shape of datatypes and constructors is managed in CakeML's environment. This particular piece of information is allowed to vary in closures, since ML supports local type definitions. Tracking this would greatly complicate our proofs. Hence, we fix a global set of constructors and enforce that all values use exactly that one.

5.8.4 Correspondence relation. We define two different correspondence relations: One for values and one for expressions.

Definition 5.20 (Expression correspondence).

$$\textsc{Var}\ \frac{}{\text{rel_e (Svar } n\text{) (Cake.Var } n\text{)}} \qquad \textsc{Const}\ \frac{n \notin \textit{constructors}}{\text{rel_e (Sconst } n\text{) (Cake.Var } n\text{)}}$$

$$\textsc{Constr}\ \frac{n \in \textit{constructors} \qquad \text{rel_e}\ t_1\ u_1 \qquad \cdots}{\text{rel_e (Sconst } \textit{name}\ \$\ t_1\ \$ \ldots \$\ t_n\text{) (Cake.Con (Some (Cake.Short } \textit{name}\text{) } [u_1, \ldots, u_n]))}$$

$$\textsc{App}\ \frac{\text{rel_e}\ t_1\ u_1 \qquad \text{rel_e}\ t_2\ u_2}{\text{rel_e}\ t_1\ \$\ t_2\ \text{Cake.App Cake.Opapp}\ [u_1, u_2]}$$

$$\textsc{Fun}\ \frac{n \notin \text{ids}\ (\Lambda\ [p_1 \Rightarrow t_1, \ldots]) \cup \textit{constructors} \qquad q_1 = \text{mk_ml_pat}\ p_1 \qquad \text{rel_e}\ t_1\ u_1 \qquad \cdots}{\text{rel_e}\ (\Lambda\ [p_1 \Rightarrow t_1, \ldots])\ (\text{Cake.Fun}\ n\ (\text{Cake.Mat (Cake.Var}\ n))\ [q_1 \Rightarrow u_1, \ldots])}$$

$$\textsc{Mat}\ \frac{\text{rel_e}\ t\ u \qquad q_1 = \text{mk_ml_pat}\ p_1 \qquad \text{rel_e}\ t_1\ u_1 \qquad \cdots}{\text{rel_e}\ (\Lambda\ [p_1 \Rightarrow t_1, \ldots]\ \$\ t)\ (\text{Cake.Mat}\ u\ [q_1 \Rightarrow u_1, \ldots])}$$

We will explain each of the rules briefly here.

Var Variables are directly related by identical name.

CONST As described earlier, constructors are treated specially in CakeML. In order to not confuse functions or variables with data constructors themselves, we require that the constant name is not a constructor.

CONSTR Constructors are directly related by identical name, and recursively related arguments.

APP CakeML does not just support general function application but also unary and binary operators. In fact, function application is the binary operator Opapp. We never generate other operators. Hence the correspondence is restricted to Opapp.

FUN/MAT Observe the symmetry between these two cases: In our term language, matching and abstraction are combined, which is not the case in CakeML. This means we relate a case abstraction to a CakeML function containing a match, and a case abstraction applied to a value to just a CakeML match.

There is no separate relation for patterns, because their translation is simple.

The value correspondence (rel_v) is structurally simpler. In the case of constructor values (Vconstr and Cake.Conv), arguments are compared recursively. Closures and recursive closures are compared extensionally, i.e. only bindings that occur in the body are checked recursively for correspondence.

5.8.5 Correctness. We use the same trick as in §5.6.4 to obtain a suitable environment for CakeML evaluation based on the rule set *rs*.

THEOREM 5.21 (CORRECTNESS). *If the compiled expression* sterm_to_cake t *terminates with a value* u *in the CakeML semantics, there is a value* v *such that* rel_v v u *and* $rs \vdash t \downarrow v$.

6 COMPOSITION

The complete compiler pipeline consists of multiple phases. Correctness is justified for each phase between intermediate semantics and correspondence relations, most of which are rather technical. Whereas the compiler may be complex and impenetrable, the trustworthiness of the constructions hinges on the obviousness of those correspondence relations.

Fortunately, under the assumption that terms to be evaluated and the resulting values do not contain abstractions – or closures, respectively – all of the correspondence relations collapse to simple structural equality: two terms are related if and only if one can be converted to the other by consistent renaming of term constructors.

For example, to convert a Cake.v back into a value:

fun cake_to_value :: Cake.v ⇒ value **where**
cake_to_value (Cake.Conv (Some (*name*, _)) *vs*) = Vconstr *name* (map cake_to_value *vs*)

This function is intentionally underspecified, because we assume that it is not applicable for closures.

The actual compiler can be characterized with two functions. Firstly, the translation of term to Cake.exp is a simple composition of each term translation function:

definition term_to_cake :: term ⇒ Cake.exp **where**
term_to_cake = sterm_to_cake ∘ pterm_to_sterm ∘ nterm_to_pterm ∘ term_to_nterm

Secondly, the function that translates function definitions by composing the phases as outlined in Figure 3, including iterated application of pattern elimination:

definition compile :: (term × term) fset ⇒ Cake.dec **where**
compile = Cake.Dletrec ∘ compile_srules_to_cake ∘ compile_prules_to_srules ∘ compile_irules_to_srules ∘
compile_irules_iter ∘ compile_crules_to_irules ∘ consts_of ∘ compile_rules_to_nrules

Each function compile_* corresponds to one compiler phase; the remaining functions are trivial. This produces a CakeML top-level declaration. We prove that evaluating this declaration in the top-level semantics (evaluate_prog) results in an environment cake_sem_env. But cake_sem_env can also be computed via another instance of the global clause set trick (§5.6.4).

Equipped with these functions, we can state the final correctness theorem:

theorem compiled_correct:
  (If CakeML evaluation of a term succeeds ... *)*
  assumes evaluate False cake_sem_env s (term_to_cake t) (s', Rval ml_v)
  (... producing a constructor term without closures ... *)*
  assumes cake_abstraction_free ml_v
  (... and some syntactic properties of the involved terms hold ... *)*
  assumes closed t **and** ¬ shadows_consts (heads $rs \cup$ *constructors*) t **and**
    welldefined (heads $rs \cup$ *constructors*) t **and** wellformed t
  (... then this evaluation can be reproduced in the term–rewriting semantics *)*
  shows $rs \vdash t \longrightarrow^*$ cake_to_term ml_v

This theorem directly relates the evaluation of a term t in the full CakeML (including mutability and exceptions) to the evaluation in the initial higher-order term rewriting semantics. The evaluation of t happens using the environment produced from the initial rule set. Hence, the theorem can be interpreted as the correctness of the pseudo-ML expression **let rec** rs **in** t.

Observe that in the assumption, the conversion goes from our terms to CakeML expressions, whereas in the conclusion, the conversion goes the opposite direction.

7 DICTIONARY CONSTRUCTION

Isabelle's type system supports *type classes* (or simply *classes*) [18, 42] whereas CakeML does not. In order to not complicate the correctness proofs, type classes are not supported by our embedded term language either. Instead, we eliminate classes and instances by a dictionary construction [19] before embedding into the term language. Haftmann and Nipkow give a pen-and-paper correctness proof of this construction [17, §4.1]. We augmented the dictionary construction with the generation of a certificate theorem that shows the equivalence of the two versions of a function, with type classes and with dictionaries. This section briefly explains our dictionary construction.

Figure 12 shows a simple example of a dictionary construction. Type variables may carry *class constraints* (e.g. α :: add). The basic idea is that classes become *dictionaries* containing the functions of that class; class instances become dictionary definitions. Dictionaries are realized as datatypes. Class constraints become additional dictionary parameters for that class. In the example, class add becomes dict_add; function f is translated into f' which takes an additional parameter of type dict_add. In reality our tool does not produce the Isabelle source code shown in Figure 12 (b) but performs the constructions internally. The correctness lemma f′_eq is proved automatically. Its precondition expresses that the dictionary must contain exactly the function(s) of class add. For any monomorphic instance, the precondition can be proved outright based on the certificate theorems proved for each class instance as explained next.

Not shown in the example is the translation of class instances. The basic form of a class instance in Isabelle is $\tau :: (c_1, \ldots, c_n)\ c$ where τ is an n-ary type constructor. It corresponds to Haskell's $(c_1\,\alpha_1, \ldots, c_n\,\alpha_n) \Rightarrow c\,(\tau\,\alpha_1 \ldots \alpha_n)$ and is translated into a function $\mathsf{inst_}c\mathsf{_}\tau :: \alpha_1\,\mathsf{dict_}c_1 \Rightarrow \cdots \Rightarrow \alpha_n\,\mathsf{dict_}c_n \Rightarrow (\alpha_1, \ldots, \alpha_n)\,\tau\,\mathsf{dict_}c$ and the following certificate theorem is proved:

$$\mathsf{cert_}c_1\ dict_1 \longrightarrow \cdots \longrightarrow \mathsf{cert_}c_n\ dict_n \longrightarrow \mathsf{cert_}c\ (\mathsf{inst_}c\mathsf{_}\tau\ dict_1\ \ldots\ dict_n)$$

```
class add =
  fixes plus :: 'a ⇒ 'a ⇒ 'a

definition f :: ('a::add) ⇒ 'a where
f x = plus x x
```

(a) Source program

```
datatype 'a dict_add = Dict_add ('a ⇒ 'a ⇒ 'a)

fun cert_add :: ('a::add) dict_add ⇒ bool where
cert_add (Dict_add pls) = (pls = plus)

fun f' :: 'a dict_add ⇒ 'a ⇒ 'a where
f' (Dict_add pls) x = pls x x

lemma f'_eq: cert_add dict ⟶ f' dict = f
<proof>
```

(b) Result of translation

Fig. 12. Dictionary construction in Isabelle

8 CONCLUSION

For this paper we have concentrated on the compiler from Isabelle/HOL to CakeML abstract syntax trees. Partial correctness is proved w.r.t. the big-step semantics of CakeML. In the next step we will link our work with the compiler from CakeML to machine code. Tan *et al.* [38, §10] prove a correctness theorem that relates their semantics with the execution of the compiled machine code. In that paper, they use a newer iteration of the CakeML semantics (functional big-step [33]) than we do here. Both semantics are still present in the CakeML source repository, together with an equivalence proof.

Evaluation of our compiled programs is already possible via Isabelle's predicate compiler [6], which allows us to turn CakeML's big-step semantics into an executable function. We have used this execution mechanism to establish for sample programs that they terminate successfully. We also plan to prove that our compiled programs terminate, i.e. total correctness.

The total size of this formalization, excluding theories extracted from Lem, is currently approximately 20000 lines of proof text (90 %) and ML code (10 %).

REFERENCES

[1] 2014. *The HOL System Description.* https://hol-theorem-prover.org/

[2] Abhishek Anand, Andrew W Appel, Greg Morrisett, Zoe Paraskevopoulou, Randy Pollack, Olivier Savary Bélanger, Matthieu Sozeau, and Matthew Weaver. 2017. CertiCoq: A verified compiler for Coq. In *CoqPL'17: The Third International Workshop on Coq for Programming Languages.* http://www.cs.princeton.edu/~appel/papers/certicoq-coqpl.pdf

[3] Lennart Augustsson. 1985. Compiling pattern matching. In *Functional Programming Languages and Computer Architecture.* Springer Berlin Heidelberg, 368–381. DOI:http://dx.doi.org/10.1007/3-540-15975-4_48

[4] Clemens Ballarin. 2016. *Tutorial to Locales and Locale Interpretation.* https://isabelle.in.tum.de/dist/Isabelle2016-1/doc/locales.pdf

[5] Nick Benton and Chung-Kil Hur. 2009. Biorthogonality, step-indexing and compiler correctness. In *Proceeding of the 14th ACM SIGPLAN international conference on Functional programming, ICFP 2009*, Graham Hutton and Andrew P. Tolmach (Eds.). ACM, 97–108. DOI:http://dx.doi.org/10.1145/1596550.1596567

[6] Stefan Berghofer, Lukas Bulwahn, and Florian Haftmann. 2009. Turning Inductive into Equational Specifications. In *Lecture Notes in Computer Science.* Springer Berlin Heidelberg, 131–146. DOI:http://dx.doi.org/10.1007/978-3-642-03359-9_11

[7] Stefan Berghofer and Tobias Nipkow. 2002. Executing Higher Order Logic. In *Types for Proofs and Programs (TYPES 2000)*, P. Callaghan, Z. Luo, J. McKinna, and R. Pollack (Eds.), Vol. 2277. 24–40.

[8] Jasmin Christian Blanchette, Johannes Hölzl, Andreas Lochbihler, Lorenz Panny, Andrei Popescu, and Dmitriy Traytel. 2014. Truly Modular (Co)datatypes for Isabelle/HOL. In *Interactive Theorem Proving.* Springer International Publishing, 93–110. DOI:http://dx.doi.org/10.1007/978-3-319-08970-6_7

[9] Mathieu Boespflug, Maxime Dénès, and Benjamin Grégoire. 2011. Full Reduction at Full Throttle. In *Certified Programs and Proofs (Lecture Notes in Computer Science)*, Jean-Pierre Jouannaud and Zhong Shao (Eds.), Vol. 7086. Springer, 362–377.

[10] Robert S. Boyer and J. Strother Moore. 2002. Single-Threaded Objects in ACL2. In *Practical Aspects of Declarative Languages (PADL 2002) (Lecture Notes in Computer Science)*, Shriram Krishnamurthi and C. R. Ramakrishnan (Eds.), Vol. 2257. Springer, 9–27. DOI:http://dx.doi.org/10.1007/3-540-45587-6_3

[11] Adam Chlipala. 2010. A verified compiler for an impure functional language. In *Proceedings of the 37th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, POPL 2010*, Manuel V. Hermenegildo and Jens Palsberg (Eds.). ACM, 93–106. DOI:http://dx.doi.org/10.1145/1706299.1706312

[12] Judy Crow, Sam Owre, John Rushby, N. Shankar, and Dave Stringer-Calvert. 2001. *Evaluating, Testing, and Animating PVS Specifications.* Technical Report. Computer Science Laboratory, SRI International, Menlo Park, CA.

[13] N. G. de Bruijn. 1972. Lambda calculus notation with nameless dummies, a tool for automatic formula manipulation, with application to the Church-Rosser theorem. *Indagationes Mathematicae (Proceedings)* 75, 5 (1972), 381 – 392. DOI:http://dx.doi.org/10.1016/1385-7258(72)90034-0

[14] Arthur D. Flatau. 1992. *A Verified Implementation of an Applicative Language with Dynamic Storage Allocation.* Ph.D. Dissertation. University of Texas at Austin.

[15] Yannick Forster and Fabian Kunze. 2016. Verified Extraction from Coq to a Lambda-Calculus. In *The 8th Coq Workshop.* http://www.ps.uni-saarland.de/~forster/coq-workshop-16/abstract-coq-ws-16.pdf

[16] David A. Greve, Matt Kaufmann, Panagiotis Manolios, J. Strother Moore, Sandip Ray, José-Luis Ruiz-Reina, Rob Sumners, Daron Vroon, and Matthew Wilding. 2008. Efficient execution in an automated reasoning environment. *J. Funct. Program.* 18, 1 (2008), 15–46. DOI:http://dx.doi.org/10.1017/S0956796807006338

[17] Florian Haftmann and Tobias Nipkow. 2010. Code Generation via Higher-Order Rewrite Systems. In *Functional and Logic Programming: 10th International Symposium, FLOPS 2010, Sendai, Japan, April 19-21, 2010. Proceedings*, Matthias Blume, Naoki Kobayashi, and Germán Vidal (Eds.). Springer Berlin Heidelberg, Berlin, Heidelberg, 103–117. DOI:http://dx.doi.org/10.1007/978-3-642-12251-4_9

[18] Florian Haftmann and Makarius Wenzel. 2007. Constructive Type Classes in Isabelle. In *Types for Proofs and Programs: International Workshop, TYPES 2006, Nottingham, UK, April 18-21, 2006, Revised Selected Papers*, Thorsten Altenkirch and Conor McBride (Eds.). Springer Berlin Heidelberg, Berlin, Heidelberg, 160–174. DOI:http://dx.doi.org/10.1007/978-3-540-74464-1_11

[19] Cordelia V. Hall, Kevin Hammond, Simon L. Peyton Jones, and Philip L. Wadler. 1996. Type classes in Haskell. *ACM Transactions on Programming Languages and Systems* 18, 2 (mar 1996), 109–138. DOI:http://dx.doi.org/10.1145/227699.227700

[20] Claudio Hermida, Uday S. Reddy, and Edmund P. Robinson. 2014. Logical Relations and Parametricity – A Reynolds Programme for Category Theory and Programming Languages. *Electronic Notes in Theoretical Computer Science* 303 (2014), 149 – 180. DOI:http://dx.doi.org/10.1016/j.entcs.2014.02.008

[21] Alexander Krauss. 2010. Partial and Nested Recursive Function Definitions in Higher-order Logic. *Journal of Automated Reasoning* 44, 4 (2010), 303–336. DOI:http://dx.doi.org/10.1007/s10817-009-9157-2

[22] Alexander Krauss and Andreas Schropp. 2010. A Mechanized Translation from Higher-Order Logic to Set Theory. In *Interactive Theorem Proving.* Springer Berlin Heidelberg, 323–338. DOI:http://dx.doi.org/10.1007/978-3-642-14052-5_23

[23] Ramana Kumar, Magnus O. Myreen, Michael Norrish, and Scott Owens. 2014. CakeML: A Verified Implementation of ML. In *Proceedings of the 41st ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages (POPL '14).* ACM, New York, NY, USA, 179–191. DOI:http://dx.doi.org/10.1145/2535838.2535841

[24] P. J. Landin. 1964. The Mechanical Evaluation of Expressions. *Comput. J.* 6, 4 (jan 1964), 308–320. DOI:http://dx.doi.org/10.1093/comjnl/6.4.308

[25] Xavier Leroy. 2009. Formal Verification of a Realistic Compiler. *Commun. ACM* 52, 7 (July 2009), 107–115. DOI:http://dx.doi.org/10.1145/1538788.1538814

[26] Pierre Letouzey. 2003. A New Extraction for Coq. In *Types for Proofs and Programs (Lecture Notes in Computer Science)*, Herman Geuvers and Freek Wiedijk (Eds.), Vol. 2646. Springer, 200–219.

[27] Robin Milner, Mads Tofte, Robert Harper, and David MacQueen. 1997. *The definition of Standard ML (revised).* MIT Press.

[28] Dominic P. Mulligan, Scott Owens, Kathryn E. Gray, Tom Ridge, and Peter Sewell. 2014. Lem: Reusable Engineering of Real-world Semantics. In *Proceedings of the 19th ACM SIGPLAN International Conference on Functional Programming (ICFP '14).* ACM, New York, NY, USA, 175–188. DOI:http://dx.doi.org/10.1145/2628136.2628143

[29] Magnus O. Myreen and Scott Owens. 2014. Proof-producing translation of higher-order logic into pure and stateful ML. *Journal of Functional Programming* 24, 2-3 (001 005 2014), 284–315. DOI:http://dx.doi.org/10.1017/S0956796813000282

[30] Georg Neis, Chung-Kil Hur, Jan-Oliver Kaiser, Craig McLaughlin, Derek Dreyer, and Viktor Vafeiadis. 2015. Pilsner: A Compositionally Verified Compiler for a Higher-order Imperative Language. In *Proceedings of the 20th ACM SIGPLAN International Conference on Functional Programming (ICFP 2015)*. ACM, New York, NY, USA, 166–178. DOI:http://dx.doi.org/10.1145/2784731.2784764

[31] Tobias Nipkow and Gerwin Klein. 2014. *Concrete Semantics*. Springer. http://www.concrete-semantics.org/

[32] Tobias Nipkow, Lawrence Paulson, and Markus Wenzel. 2002. *Isabelle/HOL — A Proof Assistant for Higher-Order Logic*. LNCS, Vol. 2283. 218 pp.

[33] Scott Owens, Magnus O. Myreen, Ramana Kumar, and Yong Kiam Tan. 2016. *Functional Big-Step Semantics*. Springer Berlin Heidelberg, Berlin, Heidelberg, 589–615. DOI:http://dx.doi.org/10.1007/978-3-662-49498-1_23

[34] Simon L. Peyton Jones. 1987. *The Implementation of Functional Programming Languages*. Prentice-Hall, Inc., Upper Saddle River, NJ, USA.

[35] Manfred Schmidt-Schauß and Jörg Siekmann. 1988. *Unification Algebras: An Axiomatic Approach to Unification, Equation Solving and Constraint Solving*. Technical Report SEKI-report SR-88-09. FB Informatik, Universität Kaiserslautern. http://www.ki.informatik.uni-frankfurt.de/papers/schauss/unif-algebr.pdf

[36] Natarajan Shankar. 2001. Static Analysis for Safe Destructive Updates in a Functional Language. In *Logic Based Program Synthesis and Transformation (LOPSTR 2001) (Lecture Notes in Computer Science)*, Alberto Pettorossi (Ed.), Vol. 2372. Springer, 1–24.

[37] Konrad Slind. 1999. *Reasoning about Terminating Functional Programs*. Ph.D. Dissertation. Technische Universität München. http://nbn-resolving.de/urn/resolver.pl?urn:nbn:de:bvb:91-diss1999111516455

[38] Yong Kiam Tan, Magnus O. Myreen, Ramana Kumar, Anthony Fox, Scott Owens, and Michael Norrish. 2016. A new verified compiler backend for CakeML. In *Proceedings of the 21st ACM SIGPLAN International Conference on Functional Programming - ICFP 2016*. Association for Computing Machinery (ACM). DOI:http://dx.doi.org/10.1145/2951913.2951924

[39] D. A. Turner. 2013. Some History of Functional Programming Languages. In *Lecture Notes in Computer Science*. Springer Berlin Heidelberg, 1–20. DOI:http://dx.doi.org/10.1007/978-3-642-40447-4_1

[40] Christian Urban. 2008. Nominal Techniques in Isabelle/HOL. *Journal of Automated Reasoning* 40, 4 (2008), 327–356. DOI:http://dx.doi.org/10.1007/s10817-008-9097-2

[41] Christian Urban, Stefan Berghofer, and Cezary Kaliszyk. 2013. Nominal 2. *Archive of Formal Proofs* (Feb. 2013). http://isa-afp.org/entries/Nominal2.shtml, Formal proof development.

[42] Markus Wenzel. 1997. Type classes and overloading in higher-order logic. In *Theorem Proving in Higher Order Logics: 10th International Conference, TPHOLs '97 Murray Hill, NJ, USA, August 19–22, 1997 Proceedings*, Elsa L. Gunter and Amy Felty (Eds.). Springer Berlin Heidelberg, Berlin, Heidelberg, 307–322. DOI:http://dx.doi.org/10.1007/BFb0028402