

Verified Characteristic Formulae for CakeML

Armaël Guéneau, Magnus O. Myreen, Ramana Kumar, Michael Norrish

April 27, 2017

Goal: write programs in a **high-level (ML-style) language**, **prove** them **correct interactively**, and compile them using a **fully verified compilation chain**.

Main building blocks:

- The CakeML compiler (POPL'14, ICFP'16)
- Characteristic Formulae for ML (ICFP'11)




CAKEML


A Verified Implementation of ML

- Features: references, modules, datatypes, exceptions, a FFI, ...
- Missing features: functors, module nesting, records

The CakeML compiler:

- Optimizing compiler
- Verified compiler: “Compcert for ML”
- Small trusted base: HOL kernel & machine axiomatisation
- Bootstraps (compiles itself)






Characteristic Formulae for ML (CFML)

CFML:

- A tool to reason about ML programs. . .
- . . . using Separation Logic
- . . . in an interactive proof assistant (Coq).
- Smoothly integrate the SL reasoning rules into the higher-order logic, by turning all program variables into logical variables in one step
- Used to verify numerous non-trivial data-structures and algorithms (Union-Find, Dijkstra, Binary trees, Vectors, Hashtables...)

Characteristic Formulae for ML (CFML)



Connecting CakeML and Characteristic Formulae
What are the missing bits?

End-to-end verification of ML programs: the missing bits

Main challenge: realize CF axioms against the source code semantics

Other challenges:

- Add support for exceptions
- Add support for I/O
- Adapt from Coq to HOL
- Adapt CakeML translator to integrate into CF the connection between program values and logical values

A program logic for CakeML

- State modular specifications about ML programs
- Prove them in a convenient way (using Separation Logic, following their CF)
- Theorem: the toplevel specification carries to the machine code produced by the CakeML compiler

Trusted codebase: HOL kernel, machine axiomatization

Background on CF

Soundness theorem: connecting CF to CakeML semantics

Extensions of CF


- Support for I/O through the CakeML FFI

- Support for exceptions

Interoperating with the proof-producing translator

Background on CF

How does the CF framework work?




Main workhorse: the CF generator, "cf".

- Source-level expression $e \rightarrow$ characteristic formula $(cf\ e)$

$(cf\ e)$:

- logical formula; doesn't mention the syntax of e
- akin to a total correctness Hoare triple

How does the CF framework work? (2)



(cf e) env H Q :

- “ e can have H as pre-condition and Q as post-condition in environment env ”
- H, Q : heap predicates (Separation Logic assertions)
- $H : \text{heap} \rightarrow \text{bool}$
- $Q : v \rightarrow \text{heap} \rightarrow \text{bool}$

Example: “let $x = e_1$ in e_2 ”

Hoare-logic rule:

$$\frac{env \vdash \{H\} e_1 \{Q'\} \quad \forall X. ((x, X) :: env) \vdash \{Q' X\} e_2 \{Q\}}{env \vdash \{H\} (\text{Let } x e_1 e_2) \{Q\}}$$

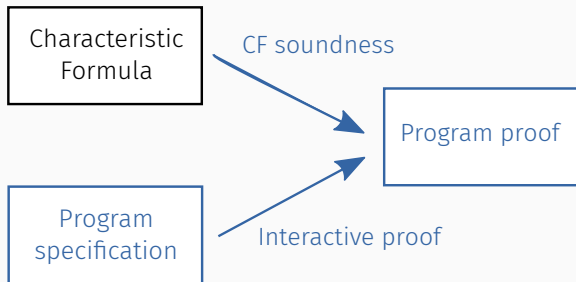
Characteristic Formula:

$$\begin{aligned} \text{cf } (\text{Let } x e_1 e_2) env = & \text{ local } (\lambda H Q. \\ & \exists Q'. \\ & \text{cf } e_1 H Q' \wedge \\ & \forall X. \text{cf } e_2 ((x, X) :: env) (Q' X) Q) \end{aligned}$$

Note: in practice, tactics are provided and the definition of `cf` is not shown to the user

Soundness theorem: connecting CF to CakeML semantics

Soundness of a CF framework



“Proving properties about a characteristic formula gives equivalent properties about the program itself”

Connecting CF and CakeML views of the heap

CakeML view of the heap: list of store values.

```
state =  
  <| clock : num  
    ; refs : v store_v list  
    ; ... |>
```

CF view of the heap: Separation Logic heap assertions.

Example: $(r_1 \rightsquigarrow v_1 * r_2 \rightsquigarrow v_2)$

```
'a store_v =  
  Refv of 'a  
  | W8array of word8 list  
  | Varray of 'a list
```

Define heaps and heap predicates specialized for CakeML values:

- Type heap = $(\text{num} \mapsto \text{v store_v})$
- Projection `state_to_heap` : `state` \rightarrow heap
- $r \rightsquigarrow v = (\lambda h. \exists loc. r = \text{Loc } loc \wedge h = \{ (loc, \text{Refv } v) \})$

Connecting logical values to CakeML deep-embedded values

CakeML values:

```
v =  
  Litv lit  
| Conv ((conN × tid_or_exn) option) (v list)  
| Closure (v sem_env) string exp  
| Recclosure (v sem_env) ((string × string × exp) list) string  
| Loc num  
| Vectorv (v list)
```

We re-use CakeML *refinement invariants*:

$$\text{INT } i = (\lambda v. v = \text{Litv } (\text{IntLit } i))$$
$$\text{BOOL T} = (\lambda v. v = \text{Conv } (\text{Some } (\text{"true"}, \text{TypeId } (\text{Short } \text{"bool"})))) []]$$
$$\vdash \text{INT } i_0 v_0 \wedge \text{INT } i_1 v_1 \Rightarrow$$
$$\{\text{emp}\} \text{plus_v} \cdot [v_0; v_1] \{\lambda v. \langle \text{INT } (i_0 + i_1) v \rangle\}$$

Realizing CFML axioms: Hoare-triple semantics

Extract from CakeML big-step semantics:

```
evaluate st env [Lit l] = (st, Rval [Litv l])
evaluate st env [Var n] =
  case lookup_var_id n env of
  None => (st, Rerr (Rabort Rtype_error))
  | Some v => (st, Rval [v])
evaluate st env [Fun x e] = (st, Rval [Closure env x e])
evaluate st env [App Opapp [f; v]] =
  case evaluate st env [v; f] of
  (st', Rval [v; f]) =>
    case do_opapp [f; v] of
    None => (st', Rerr (Rabort Rtype_error))
    | Some (env', e) =>
      if st'.clock = 0 then
        (st', Rerr (Rabort Rtimeout_error))
      else evaluate (dec_clock st') env' [e]
  | res => res
...
```

```
evaluate :
  state ->
  v sem_env ->
  exp list ->
  state * (v list, v) result
```

Realizing CFML axioms: Hoare-triple semantics

Hoare-triple for an expression e in environment env : “ $env \vdash \{H\} e \{Q\}$ ”

$$\begin{aligned} env \vdash \{H\} e \{Q\} &\iff \\ \forall st \ H'. & \\ (H * H') (state_to_heap \ st) &\Rightarrow \\ \exists v \ st' \ ck. & \\ \text{evaluate } (st \text{ with clock } := ck) \ env \ [e] &= (st', Rval \ [v]) \wedge \\ (Q \ v * H' * true) (state_to_heap \ st') & \end{aligned}$$

- H' accounts for the framed heap
- ck accounts for termination
- $true$ accounts for discarded memory cells

Proving CF soundness

Bridging the gap between:

Characteristic formulae

$$\begin{aligned} \text{cf } (\text{Let } x \ e_1 \ e_2) \ \text{env} &= \text{local } (\lambda H \ Q. \\ &\exists Q'. \\ &\text{cf } e_1 \ H \ Q' \wedge \\ &\forall xv. \text{cf } e_2 \ ((x, xv) :: \text{env}) \ (Q' \ xv) \ Q) \end{aligned}$$

Big-step semantics

$$\begin{aligned} \text{evaluate } st \ \text{env} \ [\text{Let } x \ e_1 \ e_2] &= \\ \text{case evaluate } st \ \text{env} \ [e_1] \ \text{of} & \\ (st', \text{Rval } v) \Rightarrow & \\ \text{evaluate } st' \ ((x, \text{HD } v) :: \text{env}) \ [e_2] & \\ | (st', \text{Rerr } v) \Rightarrow (st', \text{Rerr } v) & \end{aligned}$$

Proving CF soundness

Theorem (CF are sound wrt. CakeML semantics):

$$\vdash (\text{cf } e) \text{ env } H Q \Rightarrow \text{env} \vdash \{H\} e \{Q\}$$

Proof: by induction on the size of e .

Corollary:

If “ $(\text{cf } e) \text{ env } H Q$ ” holds, then starting from a state satisfying H , evaluating e terminates with a value v and a new state satisfying $Q v$.

Extensions of CF

Extensions of CF

Support for I/O through the CakeML
FFI

Performing I/O in CakeML

- Needed to interact with the external world (printing to `stdout`, opening files...)
- Done by calling external (C) code

CakeML I/O semantics

- State of the “external world” modeled by the semantics FFI state (what has been printed to stdout, which files are open, ...)
- Executing an FFI operation updates the state of the FFI
- FFI state changes are modeled by an oracle function
- Modular proofs: need to be able to split the FFI state using “*” (proofs about stdout may be independent from proofs about the file-system...)

```
state =  
  <| clock : num  
    ; refs : v store_v list  
    ; ffi :  $\theta$  ffi_state  
    ; ... |>
```

```
 $\theta$  ffi_state =  
  <| oracle :  
    string  $\rightarrow$   $\theta$   $\rightarrow$  byte list  $\rightarrow$   
     $\theta$  oracle_result  
    ; ffi_state :  $\theta$   
    ; ... |>
```

Problem: we know nothing about the type variable θ !

Splitting the FFI state

Solution:

- User provides: how to split the FFI state into *independent parts* (e.g. one part about stdout, one part about the file-system...)
- Each part models a fraction of the external world
- Several external functions can update the same part
- FFI parts exposed in the heap, and are *-separated

⇒ Programs compose as long as their specifications agree on parts for external functions they both use

Example: a specification for cat

```
fun do_onefile fname =
  let
    val fd = CharIO.openIn fname
    fun recurse () =
      case CharIO.fgetc fd of
        NONE => ()
      | SOME c =>
          CharIO.write c;
          recurse ()
  in recurse ();
    CharIO.close fd
  end

fun cat fnames =
  case fnames of
    [] => ()
  | f::fs => do_onefile f; cat fs
```

```
⊢ LIST FILENAME fns fnsv ∧
  every (λ fnm. inFS_fname fnm fs) fns ∧
  numOpenFDs fs < 255 ⇒
  {CATFS fs * STDOUT out}
  cat_v · [fnsv]
  {λ u.
    ⟨UNIT () u⟩ * CATFS fs *
    STDOUT (out @ catfiles_string fs fns)}
```

Extensions of CF

Support for exceptions

Without support for exceptions:

- An expression must reduce to a value
- Post-conditions have type $v \rightarrow \text{heap} \rightarrow \text{bool}$

We now allow expressions to raise an exception:

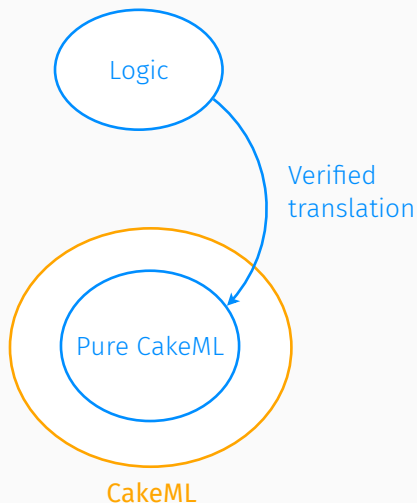
- Define datatype $\text{res} = \text{Val } v \mid \text{Exn } v$
- Post-conditions have type $\text{res} \rightarrow \text{heap} \rightarrow \text{bool}$
- Update the semantics and cf definitions accordingly...

Example: prove a more general specification for `cat`, which doesn't require that the input files exist.

Thanks to our tactics, programs without exceptions can be verified without additional effort.

Interoperating with the proof-producing translator

Existing tool: verified translation from HOL to CakeML



Produces theorems relating a (purely functional) piece of CakeML code with the corresponding HOL (pure) function.

Used to verify most of the compiler

Example: verified translation of length

- Define and verify the program in HOL4:

$$\begin{aligned} &(\text{length } [] = 0) \wedge \\ &(\text{length } (h :: t) = 1 + \text{length } t) \end{aligned}$$

$$\vdash \forall x y. \text{length } (x ++ y) = \text{length } x + \text{length } y$$

- The translator automatically produces CakeML code ...

```
fun length_ml x =
  case x of
  | [] => 0
  | (h::t) => 1 + length t
```

- ...and the theorems

```
 $\vdash$  run_prog length_ml length_env
 $\vdash$  lookup_var "length_ml" length_env = Some length_v
 $\vdash$  (a LIST  $\longrightarrow$  NUM) length length_v
```

Relating translator specifications and CF specifications

We prove equivalence between translator specifications and a particular shape of CF specifications:

$$\vdash (a \longrightarrow b) f fv \iff \forall x xv. a x xv \Rightarrow \{\text{emp}\} fv \cdot xv \{\lambda v. \langle b (f x) v \rangle\}$$

Makes translated programs and **imperative** code **with empty input and output heap** interchangeable.

Interest:

- Get specifications “for free” for pure (translated) functions
- Allows for efficient imperative implementations of algorithms that realize logical functions



- End-to-end verification framework for CakeML
- With a CF generator that supports all the features of CakeML
- With support for formally relating a logical function with an efficient imperative implementation

Future work:

- Enhance tactics to further improve proof automation
- Prove correct more data-structures and algorithms (e.g. replace bits of the CakeML compiler with more efficient, imperative code)

$$\begin{aligned}
\text{cf } (\text{Var } x) \text{ env} &= \text{local } (\lambda H Q. \\
&\quad \exists X. \text{lookup_var_id } x \text{ env} = \text{Some } X \wedge \\
&\quad H \triangleright Q X) \\
\text{cf } (\text{Let } x \ e_1 \ e_2) \text{ env} &= \text{local } (\lambda H Q. \\
&\quad \exists Q'. \\
&\quad \text{cf } e_1 \ H \ Q' \wedge \\
&\quad \forall X. \text{cf } e_2 \ ((x, X) :: \text{env}) \ (Q' \ X) \ Q) \\
\text{cf } (\text{If } \textit{cond} \ e_1 \ e_2) \text{ env} &= \text{local } (\lambda H Q. \\
&\quad \exists \textit{condv } b. \\
&\quad \text{exp_is_val } \textit{env } \textit{cond} = \text{Some } \textit{condv} \wedge \\
&\quad \text{BOOL } b \ \textit{condv} \wedge \\
&\quad ((b \iff \text{T}) \Rightarrow \text{cf } e_1 \ \textit{env} \ H \ Q) \wedge \\
&\quad ((b \iff \text{F}) \Rightarrow \text{cf } e_2 \ \textit{env} \ H \ Q))
\end{aligned}$$

Specifications:

- Written $\{H\} f \cdot args \{Q\}$: Hoare-triple for functional applications
- Related to `cf` via a consequence of the soundness theorem:

$$\begin{aligned} \vdash ns \neq [] &\Rightarrow \\ \text{length } xvs = \text{length } ns &\Rightarrow \\ cf \text{ body } (\text{extend_env } ns \ xvs \ env) \ H \ Q &\Rightarrow \\ \{H\} \text{naryClosure } env \ ns \ \text{body} \cdot xvs \ \{Q\} & \end{aligned}$$

Connecting CF and CakeML visions of the heap

Define heaps holding CakeML values:

$$\text{heap} = (\text{num} \times \text{v store_v}) \text{ set}$$
$$r \rightsquigarrow v = (\lambda h. \exists \text{loc}. r = \text{Loc } \text{loc} \wedge h = \{ (\text{loc}, \text{Refv } v) \})$$
$$p * q = (\lambda h. \exists u v. \text{split } h (u, v) \wedge p u \wedge q v)$$

Define $\text{state_to_heap} : \text{state} \rightarrow \text{heap}$.

For a state st with $st.\text{refs} = [\text{Refv } v_1; \text{Refv } v_2]$:

- $\text{state_to_heap } st = \{(0, v_1); (1, v_2)\}$
- $(\text{Loc } 0 \rightsquigarrow v_1 * \text{Loc } 1 \rightsquigarrow v_2) (\text{state_to_heap } st)$

Hoare-triple for the application of a closure to a single argument:

“ $\{H\} f \cdot x \{Q\}$ ”

$$\{H\} f \cdot x \{Q\} \iff$$

case do_opapp [f; x] of

- None $\Rightarrow \forall st\ h_1\ h_2. \text{split}(\text{state_to_heap } p\ st) (h_1, h_2) \Rightarrow \neg H\ h_1$
- | Some (env, exp) $\Rightarrow \text{env} \vdash \{H\} \text{exp} \{Q\}$

Semantics of Hoare-triples for n-ary application

Hoare-triple for the application of a closure to multiple arguments:

“ $\{H\} f \cdot args \{Q\}$ ”

$$\{H\} f \cdot [] \{Q\} \iff \mathbf{F}$$

$$\{H\} f \cdot [x] \{Q\} \iff \{H\} f \cdot x \{Q\}$$

$$\{H\} f \cdot x :: x' :: xs \{Q\} \iff$$

$$\{H\} f \cdot x \{\lambda g. \exists H'. H' * \langle \{H'\} g \cdot x' :: xs \{Q\} \rangle\}$$

Specifications are modular: app integrates the frame rule

$(\text{cf } e) \text{ env } H Q \Rightarrow$

$\forall st.$

$H(\text{state_to_heap } st) \Rightarrow$

$\exists v st' ck.$

$\text{evaluate } (st \text{ with clock } := ck) \text{ env } [e] = (st', \text{Rval } [v]) \wedge$

$(Q v * \text{true}) (\text{state_to_heap } st')$

Performing I/O in CakeML

CakeML programs do I/O using a byte-array-based foreign-function interface (FFI).

- “App (FFI *name*) [*array*]”: a CakeML expression
- Calls the external function “*name*” (typically implemented in C) with “*array*” as a parameter
- Reads back the result in “*array*”

For example: read a character from `stdin`, open a file, ...

Splitting the FFI state

Solution: parametrize `state_to_heap` with information on how to split the FFI state into “parts”.

- A *part* represents an independent bit of the external world
- Several external functions can update the same part
- The FFI state θ can be split into separated parts
- “`stdout`” would be a part, “`stdin`” an other, the filesystem a third one...

Splitting the FFI state (2)

We parametrize `state_to_heap` with:

- A projection function $proj : \theta \rightarrow (\text{string} \mapsto \text{ffi})$
- A list of FFI *parts* : $(\text{string list} \times \text{ffi_next}) \text{ list}$

`ffi`: low-level generic model for
the state of a FFI part

`ffi_next`: “next-state
function”, a part of the oracle

```
ffi =  
  Str string  
  | Num num  
  | Cons ffi ffi  
  | List (ffi list)  
  | Stream (num stream)  
  
ffi_next =  
  string  $\rightarrow$  byte list  $\rightarrow$  ffi  $\rightarrow$   
  (byte list  $\times$  ffi) option
```

Splitting the FFI state (3)

Finally, we define a generic IO heap assertion:

$$\begin{aligned} \text{IO} &: \text{ffi} \rightarrow \text{ffi_next} \rightarrow \text{string list} \rightarrow \text{heap} \rightarrow \text{bool} \\ \text{IO } st \ u \ ns &= (\lambda s. \exists ts. s = \{ \text{FFI_part } st \ u \ ns \ ts \}) \end{aligned}$$

Pre- and post-conditions can now make assertions about I/O. Users typically define more specialized assertions on top of IO.

Example: a more general specification for `cat1`

We can remove the precondition that the input file must exist:

$$\begin{aligned} &\vdash \text{FILENAME } fnm \text{ } fnv \wedge \text{numOpenFDs } fs < 255 \Rightarrow \\ &\quad \{ \text{CATFS } fs * \text{STDOUT } out \} \\ &\quad \text{cat1_v} \cdot [fnv] \\ &\quad \{ \text{POST} \\ &\quad \quad (\lambda u. \\ &\quad \quad \quad \exists content. \\ &\quad \quad \quad \langle \text{UNIT } () \ u \rangle * \langle \text{alist_lookup } fs.files \ fnm = \text{Some } content \rangle * \\ &\quad \quad \quad \text{CATFS } fs * \text{STDOUT } (out \ @ \ content)) \\ &\quad \quad (\lambda e. \\ &\quad \quad \quad \langle \text{BadFileName_exn } e \rangle * \langle \neg \text{inFS_fname } fnm \ fs \rangle * \text{CATFS } fs * \\ &\quad \quad \quad \text{STDOUT } out) \} \end{aligned}$$

Exception-aware Hoare-triples

Hoare-triple validity “ $env \vdash \{H\} e \{Q\}$ ” becomes:

$$\begin{aligned} env \vdash \{H\} e \{Q\} &\iff \\ \forall st \ h_i \ h_k. & \\ \text{split} (\text{state_to_heap } p \ st) (h_i, h_k) \Rightarrow & \\ H \ h_i \Rightarrow & \\ \exists r \ st' \ h_f \ h_g \ ck. & \\ \text{split3} (\text{state_to_heap } p \ st') (h_f, h_k, h_g) \wedge Q \ r \ h_f \wedge & \\ \text{case } r \ \text{of} & \\ \quad \text{Val } v \Rightarrow \text{evaluate } (st \ \text{with } \text{clock} := ck) \ env \ [e] = (st', \text{Rval } [v]) & \\ \quad | \ \text{Exn } v \Rightarrow \text{evaluate } (st \ \text{with } \text{clock} := ck) \ env \ [e] = (st', \text{Rerr } (\text{Rraise } v)) & \end{aligned}$$

Note: we still rule out actual failures, where `evaluate` returns “`Rerr (Rabort abort)`”.

Updating cf

Add side-conditions to characteristic formulae, to deal with exceptions:

$$\begin{aligned} \text{cf } p (\text{Var } name) env &= \text{local } (\lambda H Q. \\ &(\exists v. \text{lookup_var_id } name env = \text{Some } v \wedge H \triangleright Q (\text{Val } v)) \wedge \\ &Q \blacktriangleright_e \mathbf{F}) \end{aligned}$$
$$\begin{aligned} \text{cf } p (\text{Let } (\text{Some } x) e_1 e_2) env &= \text{local } (\lambda H Q. \\ &\exists Q'. \\ &\text{cf } p e_1 env H Q' \wedge Q' \blacktriangleright_e Q \wedge \\ &\forall xv. \text{cf } p e_2 ((x, xv) :: env) (Q' (\text{Val } xv)) Q) \end{aligned}$$

$$Q_1 \blacktriangleright_e Q_2 \iff \forall e. Q_1 (\text{Exn } e) \triangleright Q_2 (\text{Exn } e)$$

CFs for raise and handle

Define cf for Raise and Handle: similar to the Var and Let cases

```
cf p (Raise e) env = local (λ H Q.  
  ∃ v. exp_is_val env e = Some v ∧ H ▷ Q (Exn v) ∧ Q ▶v F)
```

```
cf p (Handle e rows) env = local (λ H Q.  
  ∃ Q'.  
    cf p e env H Q' ∧ Q' ▶v Q ∧  
    ∀ ev.  
      cf_cases ev ev (map (I ## cf p) rows) env (Q' (Exn ev)) Q)
```

$$Q_1 \blacktriangleright_v Q_2 \iff \forall e. Q_1 (\text{Val } e) \triangleright Q_2 (\text{Val } e)$$