



# CAKEML

A Verified Implementation of ML

Anthony Fox



Hugo Férée



Armaël Guéneau



Ramana Kumar



Magnus Myreen



Michael Norrish



Scott Owens



Yong Kiam Tan



S-REPLS 6

UCL

25 May, 2017

# The CakeML Project

A functional programming language

A verified compiler

Verified applications

Theorem proving technology

# Formal Verification: Two Extremes

Full functional correctness

Rich security properties

Interactive

1,000s of lines

Not necessarily mainstream

Bug finding

Simple security properties

Automatic

1,000,000 of lines

C, Java & ASM

# The CakeML Language

Design: *“The CakeML language is designed to be both easy to program in and easy to reason about formally”*

Reality: CakeML, the language  $\cong$  Standard ML without functors

i.e. with almost everything else:

- ✓ higher-order functions
- ✓ mutual recursion and polymorphism
- ✓ datatypes and (nested) pattern matching
- ✓ references and (user-defined) exceptions
- ✓ modules, signatures, abstract types
- ✓ polymorphic/byte arrays/vectors, FFI calls
- ? right-to-left evaluation, prefers curried style

# The CakeML Language

*was originally*

Design: “*The CakeML language is designed to be both easy to program in and easy to reason about formally*”

Reality: CakeML, the language  $\cong$  Standard ML without functors

i.e. with almost everything else:

- ✓ higher-order functions
- ✓ mutual recursion and polymorphism
- ✓ datatypes and (nested) pattern matching
- ✓ references and (user-defined) exceptions
- ✓ modules, signatures, abstract types
- ✓ polymorphic/byte arrays/vectors, FFI calls
- ? right-to-left evaluation, prefers curried style

# The CakeML Language

*was originally*

Design: “*The CakeML language is designed to be both easy to program in and easy to reason about formally*”

*It is still clean, but not always simple.*


Reality: CakeML, the language  $\cong$  Standard ML without functors

i.e. with almost everything else:

- ✓ higher-order functions
- ✓ mutual recursion and polymorphism
- ✓ datatypes and (nested) pattern matching
- ✓ references and (user-defined) exceptions
- ✓ modules, signatures, abstract types
- ✓ polymorphic/byte arrays/vectors, FFI calls
- ? right-to-left evaluation, prefers curried style

# The CakeML Language

Design: *“The CakeML language ~~is~~ designed to be easy to program in and easy to reason about.”*  
*It is still clean, but not always simple.*


was originally  Significant effort wrt the semantics

Reality: CakeML, the language  $\cong$  Standard ML without functors

i.e. with almost everything else:

- ✓ higher-order functions
- ✓ mutual recursion and polymorphism
- ✓ datatypes and (nested) pattern matching
- ✓ references and (user-defined) exceptions
- ✓ modules, signatures, abstract types
- ✓ polymorphic/byte arrays/vectors, FFI calls
- ? right-to-left evaluation, prefers curried style


# The CakeML Compiler



All languages communicate with the external world via a byte-array-based foreign-function interface.

<https://cakeml.org>

# The CakeML Compiler




All languages communicate with the external world via a byte-array-based foreign-function interface.

<https://cakeml.org>

execution time relative to native code compiled OCaml (red)


Performance numbers *before* bignums were added (Nov 2016)



Which colour is what ML implementation?

execution time relative to native code compiled OCaml (red)


Performance numbers *before* bignums were added (Nov 2016)



Which colour is what ML implementation?

execution time relative to native code compiled OCaml (red)

# Performance numbers *before* bignums were added (Nov 2016)




CakeML is sometimes faster than OCaml

ocaml 4.02.3 (red)      polyc v5.6 (green)      cakeml (cyan)  
smlnj v110.78 (blue)      mlton 20100608 (magenta)


execution time relative to native code compiled OCaml (red)

# Performance numbers *after* bignums were added (Feb 2017)




# Dimensions of Compiler Verification


# Dimensions of Compiler Verification



# Dimensions of Compiler Verification




# Dimensions of Compiler Verification



# Ecosystem


*Proof-producing synthesis*



*Verified parsing*



*Verified type inference*




*Proof-producing verification-condition generation*



*Also: x86 implementation with read-eval-print-loop*

# Unix-style utilities



# Applications

## *Unix-style utilities*

- cat
- sort
- grep
- diff+patch
- bootstrapped compiler



Johannes Pohjola

## *Standard library for CakeML:*

- module: char I/O stdin/stdout
- module: reading of files
- module: reading command-line arguments
- standard modules: lists, vectors, arrays, strings, characters, etc.

# Programming in HOL

Or Isabelle/HOL or  
Coq or ...

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

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

Theorem

Induct\_on `x` THEN SRW\_TAC [] []

Proof

EVAL (`length [1;2;3]`) = ( $\vdash \text{length } [1;2;3] = 3$ )

Secure  
evaluation

# Programming in HOL

~15,000 loc

(compile conf std\_in = ...)

$\vdash \forall \text{conf } p. \text{good\_init init init}' \text{ conf} \Rightarrow$   
 $\text{sem init } p = \text{sem\_x86 init}' (\text{compile conf } p)$

Theorem

[13 IL semantics, > 100,000 lop]

Proof

EVAL (compile ... "val x = ...") =  
( $\vdash \text{compile ... "val x = ..." = 0x48,0x39 ...}$ )

Secure  
evaluation

# Programming in HOL

~15,000 loc

(compile conf std\_in = ...)

$\vdash \forall \text{conf } p. \text{good\_init init init' conf} \Rightarrow$   
 $\text{sem\_init } p \text{ (compile conf p)}$

Theorem

Slow:  
15 minutes to 2 days

[13 IL semantics, > 100,000 loc]


Proof

EVAL ('compile ... "val x = ..."') =  
( $\vdash$  compile ... "val x = ..." = 0x48,0x39 ...)

Secure  
evaluation

# For Fast Execution: HOL to CakeML

```
fun length [] = 0
  | length (h::t) = 1 + length t
```

$$\vdash \text{length } (x++y) \Downarrow v \Rightarrow$$
$$\exists v_1 v_2. \text{length } x \Downarrow v_1 \wedge \text{length } y \Downarrow v_2 \wedge v = v_1 + v_2$$


CakeML  
semantics

# CFML: Characteristic Formulae for ML

Arthur Charguéraud



“CFML can be used to verify  
Caml programs using  
the Coq proof assistant.”

Arthur’s PhD topic

We want CF for CakeML

Arthur’s student Armaël Guéneau → Chalmers visit

# What is CF?

Verification conditions for ML programs.

For standard Hoare logic:

It suffices to show:

$$P \Rightarrow wp(c, Q)$$

To prove:

$$\{ P \} c \{ Q \}$$

For Caml programs:

It suffices to show:

$$cf \ e \ H \ Q$$

To prove:

$$\text{“} \{ H \} e \{ Q \} \text{”}$$

*cf is a function similar to wp.*

*It produces a verification condition (higher-order sep. logic).*

# Weaknesses of Arthur's CFML for Caml

CFML: The `cf` function is defined in OCaml (i.e. outside of Coq)

CFML: Soundness proved mostly outside of Coq (pen and paper).

CFML: Soundness proved w.r.t. idealise semantics of OCaml.

CFML: does not support I/O or exceptions.

# Aims with CakeML CF

CFML: The cf function is defined in OCaml (i.e. outside of Coq)

CFML: Soundness proved mostly outside of Coq (pen and paper).

CFML: Soundness proved w.r.t. idealise semantics of OCaml.

CFML: does not support I/O or exceptions.

# Aims with CakeML CF

CFML: The cf function is defined in OCaml (i.e. outside of Coq)

CakeML CF: defines cf as a function in the logic

CFML: Soundness proved mostly outside of Coq (pen and paper).

CFML: Soundness proved w.r.t. idealise semantics of OCaml.

CakeML CF: soundness proved in the logic w.r.t. CakeML semantics

CFML: does not support I/O or exceptions.

CakeML CF: supports all CakeML language features  
(incl. I/O and exceptions)

---

Weakness of CakeML CF: clunkier values (deep embedding), tactics etc.

# Soundness thm

CF generated proof obligation

$\vdash \text{cf } e \text{ env } H \ Q \Rightarrow$   
 $\forall st.$

CakeML semantics state

$H \ (\text{state\_to\_set } st) \Rightarrow$   
 $\exists st' \ h_f \ h_g \ v \ ck.$

$\text{evaluate } (st \text{ with clock } := ck) \text{ env } [e] = (st', \text{Rval } [v]) \wedge$   
 $\text{split } (\text{state\_to\_set } st') \ (h_f, h_g) \wedge Q \ v \ h_f$

semantics  
of CakeML  
source

*(Version before support for exceptions was added.)*

# I/O semantics in CakeML (FFI)

*The CakeML state carries an oracle (with a type variable):*

```
 $\theta$  ffi_state =  
  <| oracle : (string  $\rightarrow$   $\theta$   $\rightarrow$  byte list  $\rightarrow$   $\theta$  oracle_result);  
    ffi_state :  $\theta$ ;  
    final_event : (final_event option);  
    io_events : (io_event list) |>  
  
final_event = Final_event string (byte list) ffi_outcome  
ffi_outcome = FFI_diverged | FFI_failed  
io_event = IO_event string ((byte  $\times$  byte) list)  
 $\theta$  oracle_result = Oracle_return  $\theta$  (byte list) | Oracle_diverge | Oracle_fail
```

# I/O semantics in CakeML (FFI)

*The CakeML state carries an oracle (with a type variable):*

```
 $\theta$  ffi_state =  
  <| oracle : (string  $\rightarrow$   $\theta$   $\rightarrow$  byte list  $\rightarrow$   $\theta$  oracle_result);  
    ffi_state :  $\theta$ ;  
    final_event : (final_event option);  
    io_events : (io_event list) |>  
  
final_event = Final_event string (byte list) ffi_outcome  
ffi_outcome = FFI_diverged | FFI_failed  
io_event = IO_event string ((byte  $\times$  byte) list)  
 $\theta$  oracle_result = Oracle_return  $\theta$  (byte list) | Oracle_diverge | Oracle_fail
```

# I/O continued

**Reminder about the soundness theorem:**

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

$\forall st.$

$H \text{ (state\_to\_set } st) \Rightarrow$

$\exists st' \ h_f \ h_g \ v \ ck.$

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

$\text{split } \text{(state\_to\_set } st') \ (h_f, h_g) \wedge Q \ v \ h_f$

**Make state\_to\_set include a partitioned image of the FFI state so that we can write:**

$(\text{IO } s_1 \ u_1 \ [n] * \text{IO } s_2 \ u_2 \ ns * \dots) \text{ (state\_to\_set } pp \ st)$

**where:**

$\text{IO } st \ u \ ns = (\lambda s. \exists ts. s = \{ \text{FFI\_part } st \ u \ ns \ ts \})$

# Spec for part of cat

$\vdash$  FILENAME  $fnm\ fnv \wedge \text{numOpenFDs } fs < 255 \Rightarrow$   
    {CATFS  $fs * \text{STDOUT } out$ }  
    cat1\_v · [ $fnv$ ]  
    {POST  
    ( $\lambda u.$   
         $\exists content.$   
         $\langle \text{UNIT } ()\ u \rangle * \langle \text{alist\_lookup } fs.\text{files } fnm = \text{Some } content \rangle * \text{CATFS } fs * \text{STDOUT } (out @ content)$   
    ( $\lambda e.$   
         $\langle \text{BadFileName\_exn } e \rangle * \langle \neg \text{inFS\_fname } fnm\ fs \rangle * \text{CATFS } fs * \text{STDOUT } out$ )}

# Bootstrapping

function in logic (**compiler**)



**parsing**

**type inference**

**compilation**





# Bootstrapping



⊢ **compiler-ML** implements **compiler**

# Bootstrapping




⊢ **compiler-ML** implements **compiler**

by evaluation  
in the logic

⊢ **compiler** (**compiler-ML**) = **compiler-x86**

# Bootstrapping



function in logic (**compiler**)



Proof-producing  
synthesis (ICFP'12)

CakeML program (**compiler-ML**)

$\vdash$  **compiler-ML** implements **compiler**


by evaluation  
in the logic

$\vdash$  **compiler** (**compiler-ML**) = **compiler-x86**

$\vdash \forall c. (\text{compiler } c) \text{ implements } c$

by compiler  
correctness

# Bootstrapping



$\vdash$  **compiler-ML** implements **compiler**

by evaluation in the logic

$\vdash$  **compiler** (**compiler-ML**) = **compiler-x86**

$\vdash \forall c. (\text{compiler } c) \text{ implements } c$

by compiler correctness

**Theorem:**  $\vdash$  **compiler-x86** implements **compiler**

# To a compiler application

```
fun main u =
  let
    val cl = Commandline.arguments ()
  in
    case compiler_x64 cl (read_all []) of
      (c, e) => (print_app_list c; print_err e)
    end
end
```

```
`cl ≠ [] ∧ EVERY validArg cl ∧ LENGTH (FLAT cl) + LENGTH cl ≤ 256 ⇒
app (p:'ffi ffi_proj) ^(fetch_v "main" st)
[Conv NONE []]
(STDOUT out * STDERR err * (STDIN inp F * COMMANDLINE cl))
(POSTv uv. &UNIT_TYPE () uv *
STDOUT (out ++ (FLAT (MAP explode
                      (append (FST(compiler_x64 (TL(MAP implode cl))
                      inp)))))) *
STDERR (err ++ explode (SND(compiler_x64 (TL(MAP implode cl)) inp))) *
(STDIN "" T * COMMANDLINE cl))` ,
```

# To a compiler application

```
fun main u =  
  let  
    val cl = Commandline.arguments ()  
  in  
    case compiler_x64 cl (read_all []) of  
      (c, e) => (print_app_list c; ...  
    end
```

Good  
command line

```
`cl ≠ [] ∧ EVERY validArg cl ∧ LENGTH (FLAT cl) + LENGTH cl ≤ 256 ⇒  
app (p:'ffi ffi_proj) ^(fetch_v "main" st)  
[Conv NONE []]  
(STDOUT out * STDERR err * (STDIN inp F * COMMANDLINE cl))  
(POSTv uv. &UNIT_TYPE () uv *  
STDOUT (out ++ (FLAT (MAP explode  
                      (append (FST(compiler_x64 (TL(MAP implode cl))  
                      inp)))))) *  
STDERR (err ++ explode (SND(compiler_x64 (TL(MAP implode cl)) inp))) *  
(STDIN "" T * COMMANDLINE cl))` ,
```

# To a compiler application

```
fun main u =  
  let  
    val cl = Commandline.arguments ()  
  in  
    case compiler_x64 cl (read_all []) of  
      (c, e) => (print_app_list c; print_err e)  
    end
```

Unit arg.

```
`cl ≠ [] ∧ EVERY (λ s. (FLAT cl) + LENGTH cl ≤ 256 ⇒  
  app (p:'ffi ffi' ^ "main" st)  
  [Conv NONE []]  
  (STDOUT out * STDERR err * (STDIN inp F * COMMANDLINE cl))  
  (POSTv uv. &UNIT_TYPE () uv *  
  STDOUT (out ++ (FLAT (MAP explode  
    (append (FST(compiler_x64 (TL(MAP implode cl))  
      inp)))))) *  
  STDERR (err ++ explode (SND(compiler_x64 (TL(MAP implode cl)) inp))) *  
  (STDIN "" T * COMMANDLINE cl))` ,
```

# To a compiler application

```
fun main u =  
  let  
    val cl = Commandline.arguments ()  
  in  
    case compiler_x64 cl (read_all []) of  
      (c, e) => (print_app_list c; print_err e)  
    end
```

```
`cl ≠ [] ∧ EVERY validArg cl ∧ LENGTH (FLAT cl) + LENGTH  
app (p:'ffi ffi_proj) ^ (fetch_v "main" st)  
[Conv NONE []]  
(STDOUT out * STDERR err * (STDIN inp F * COMMANDLINE cl))  
(POSTv uv. &UNIT_TYPE () uv *  
  STDOUT (out ++ (FLAT (MAP explode  
    (append (FST(compiler_x64 (TL(MAP implode cl))  
      inp)))))) *  
  STDERR (err ++ explode (SND(compiler_x64 (TL(MAP implode cl)) inp))) *  
(STDIN "" T * COMMANDLINE cl))`
```

Precond.

# To a compiler application


```
fun main u =  
  let  
    val cl = Commandline.arguments ()  
  in  
    case compiler_x64 cl (read_all []) of  
      (c, e) => (print_app_list c; print_err e)  
    end
```

```
`cl ≠ [] ∧ EVERY validArg cl ∧ |F| cl ≤ 256 ⇒  
app (p:'ffi ffi_proj) ^ (fetch_...  
[Conv NONE []]  
(STDOUT out * STDERR err * (STD...  
(POSTv uv. &UNIT_TYPE () uv *  
STDOUT (out ++ (FLAT (MAP explode  
      (append (FST(compiler_x64 (TL(MAP implode cl))  
        inp)))))) *  
STDERR (err ++ explode (SND(compiler_x64 (TL(MAP implode cl)) inp))) *  
(STDIN "" T * COMMANDLINE cl))`
```

# To a compiler application

```
fun main u =  
  let  
    val cl = Commandline.arguments ()  
  in  
    case compiler_x64 cl (read_all []) of  
      (c, e) => (print_app_list c; print_err e)  
    end
```

```
`cl ≠ [] ∧ EVERY validArg cl ∧ LENGTH (FLAT cl) + LENGTH cl ≤ 256 ⇒  
app (p:'ffi ffi_proj) ^ (fetch_v "main"  
 [Conv NONE []]  
 (STDOUT out * STDERR err * (STDIN inp *  
 (POSTv uv. &UNIT_TYPE () uv *  
 STDOUT (out ++ (FLAT (MAP explode  
 (append (FST(compiler_x64 (TL(MAP implode cl))  
 inp)))))) *  
 STDERR (err ++ explode (SND(compiler_x64 (TL(MAP implode cl)) inp))) *  
 (STDIN "" T * COMMANDLINE cl))`
```



Output

# To a compiler application

```
fun main u =  
  let  
    val cl = Commandline.arguments ()  
  in  
    case compiler_x64 cl (read_all []) of  
      (c, e) => (print_app_list c; print_err e)  
    end
```

```
`cl ≠ [] ∧ EVERY validArg cl ∧ LENGTH (FLAT cl) + LENGTH cl ≤ 256 ⇒  
app (p:'ffi ffi_proj) ^ (fetch_v "main" st)  
[Conv NONE []]  
(STDOUT out * STDERR err * (STDIN in * COMMANDLINE cl))  
(POSTv uv. &UNIT_TYPE () uv)  
STDOUT (out ++ (FLAT (MAP (app (TL(MAP implode cl))  
  (app (TL(MAP implode cl)) inp)))) *  
STDERR (err ++ explode (SND(compiler_x64 (TL(MAP implode cl)) inp))) *  
(STDIN "" T * COMMANDLINE cl))`
```

Error msgs.

# To a compiler application

```
fun main u =
  let
    val cl = Commandline.arguments ()
  in
    case compiler_x64 cl (read_all []) of
      (c, e) => (print_app_list c; print_err e)
    end
end
```

```
`cl ≠ [] ∧ EVERY validArg cl ∧ LENGTH (FLAT cl) + LENGTH cl ≤ 256 ⇒
app (p:'ffi ffi_proj) ^(fetch_v "main" st)
[Conv NONE []]
(STDOUT out * STDERR err * (STDIN inp F * COMMANDLINE cl))
(POSTv uv. &UNIT_TYPE () uv *
STDOUT (out ++ (FLAT (MAP explode
                      (append (FST(compiler_x64 (TL(MAP implode cl))
                      inp)))))) *
STDERR (err ++ explode (SND(compiler_x64 (TL(MAP implode cl)) inp))) *
(STDIN "" T * COMMANDLINE cl))`,
```



Upcoming tutorials at  
PLDI and ICFP

<https://cakeml.org>

Goal: Build a tool path for creating fully verified applications.

Compiler for an ML-like programming language

Mechanically verified in HOL-4

A tool to support the construction of verified systems