

LFCS seminar, Edinburgh, May 2015



CakeML

A verified implementation of ML

Magnus Myreen

Chalmers University of Technology & University of Cambridge


Joint work with Ramana Kumar, Michael Norrish, Scott Owens and many more

Motivation



CakeML

A verified implementation of ML



Why?

We wanted to know
whether it's possible.

functional correctness

Background

From my PhD (2009):

Verified Lisp interpreter
in ARM, x86 and PowerPC machine code

Collaboration with Jared Davis (2011):

Verified Lisp read-eval-print loop
in 64-bit x86 machine code, with dynamic compilation
(plus verification of an ACL2-like theorem prover)

Can we do the same for ML?

A verified implementation of ML
(plus verification of a HOL-like theorem prover?)

Other HOL4 hackers also have relevant interests...

People involved



Ramana Kumar
(Uni. Cambridge)



Michael Norrish
(NICTA, ANU)

operational **semantics**

verified **compilation** from
CakeML to bytecode

verified **type** inferencer

verified **parsing** (syntax is
compatible with SML)

verified **x86** implementations

proof-producing **code
generation** from HOL




Scott Owens
(Uni. Kent)



Magnus Myreen
(Chalmers)

Proofs in HOL4

HOL4 is a **fully expansive** theorem prover:



All proofs expand at runtime into primitive inferences in the HOL4 kernel.

The kernel implements the axioms and inference rules of higher-order logic.

Overall aim

*to make proof assistants into **trustworthy** and **practical** program development platforms*

Trustworthy code extraction:

functions in HOL (shallow embedding)



proof-producing translation [ICFP'12, JFP'14]

CakeML program (deep embedding)



verified compilation of CakeML [POPL'14]

x86-64 machine code (deep embedding)

This talk

Part 1: verified implementation of CakeML

Part 2: current status, HOL light, future

Part 1: verified implementation of CakeML

POPL'14

CakeML: A Verified Implementation of ML

Ramana Kumar^{* 1} Magnus O. Myreen^{† 1} Michael Norrish² Scott Owens³

¹ Computer Laboratory, University of Cambridge, UK

² Canberra Research Lab, NICTA, Australia[‡]

³ School of Computing, University of Kent, UK

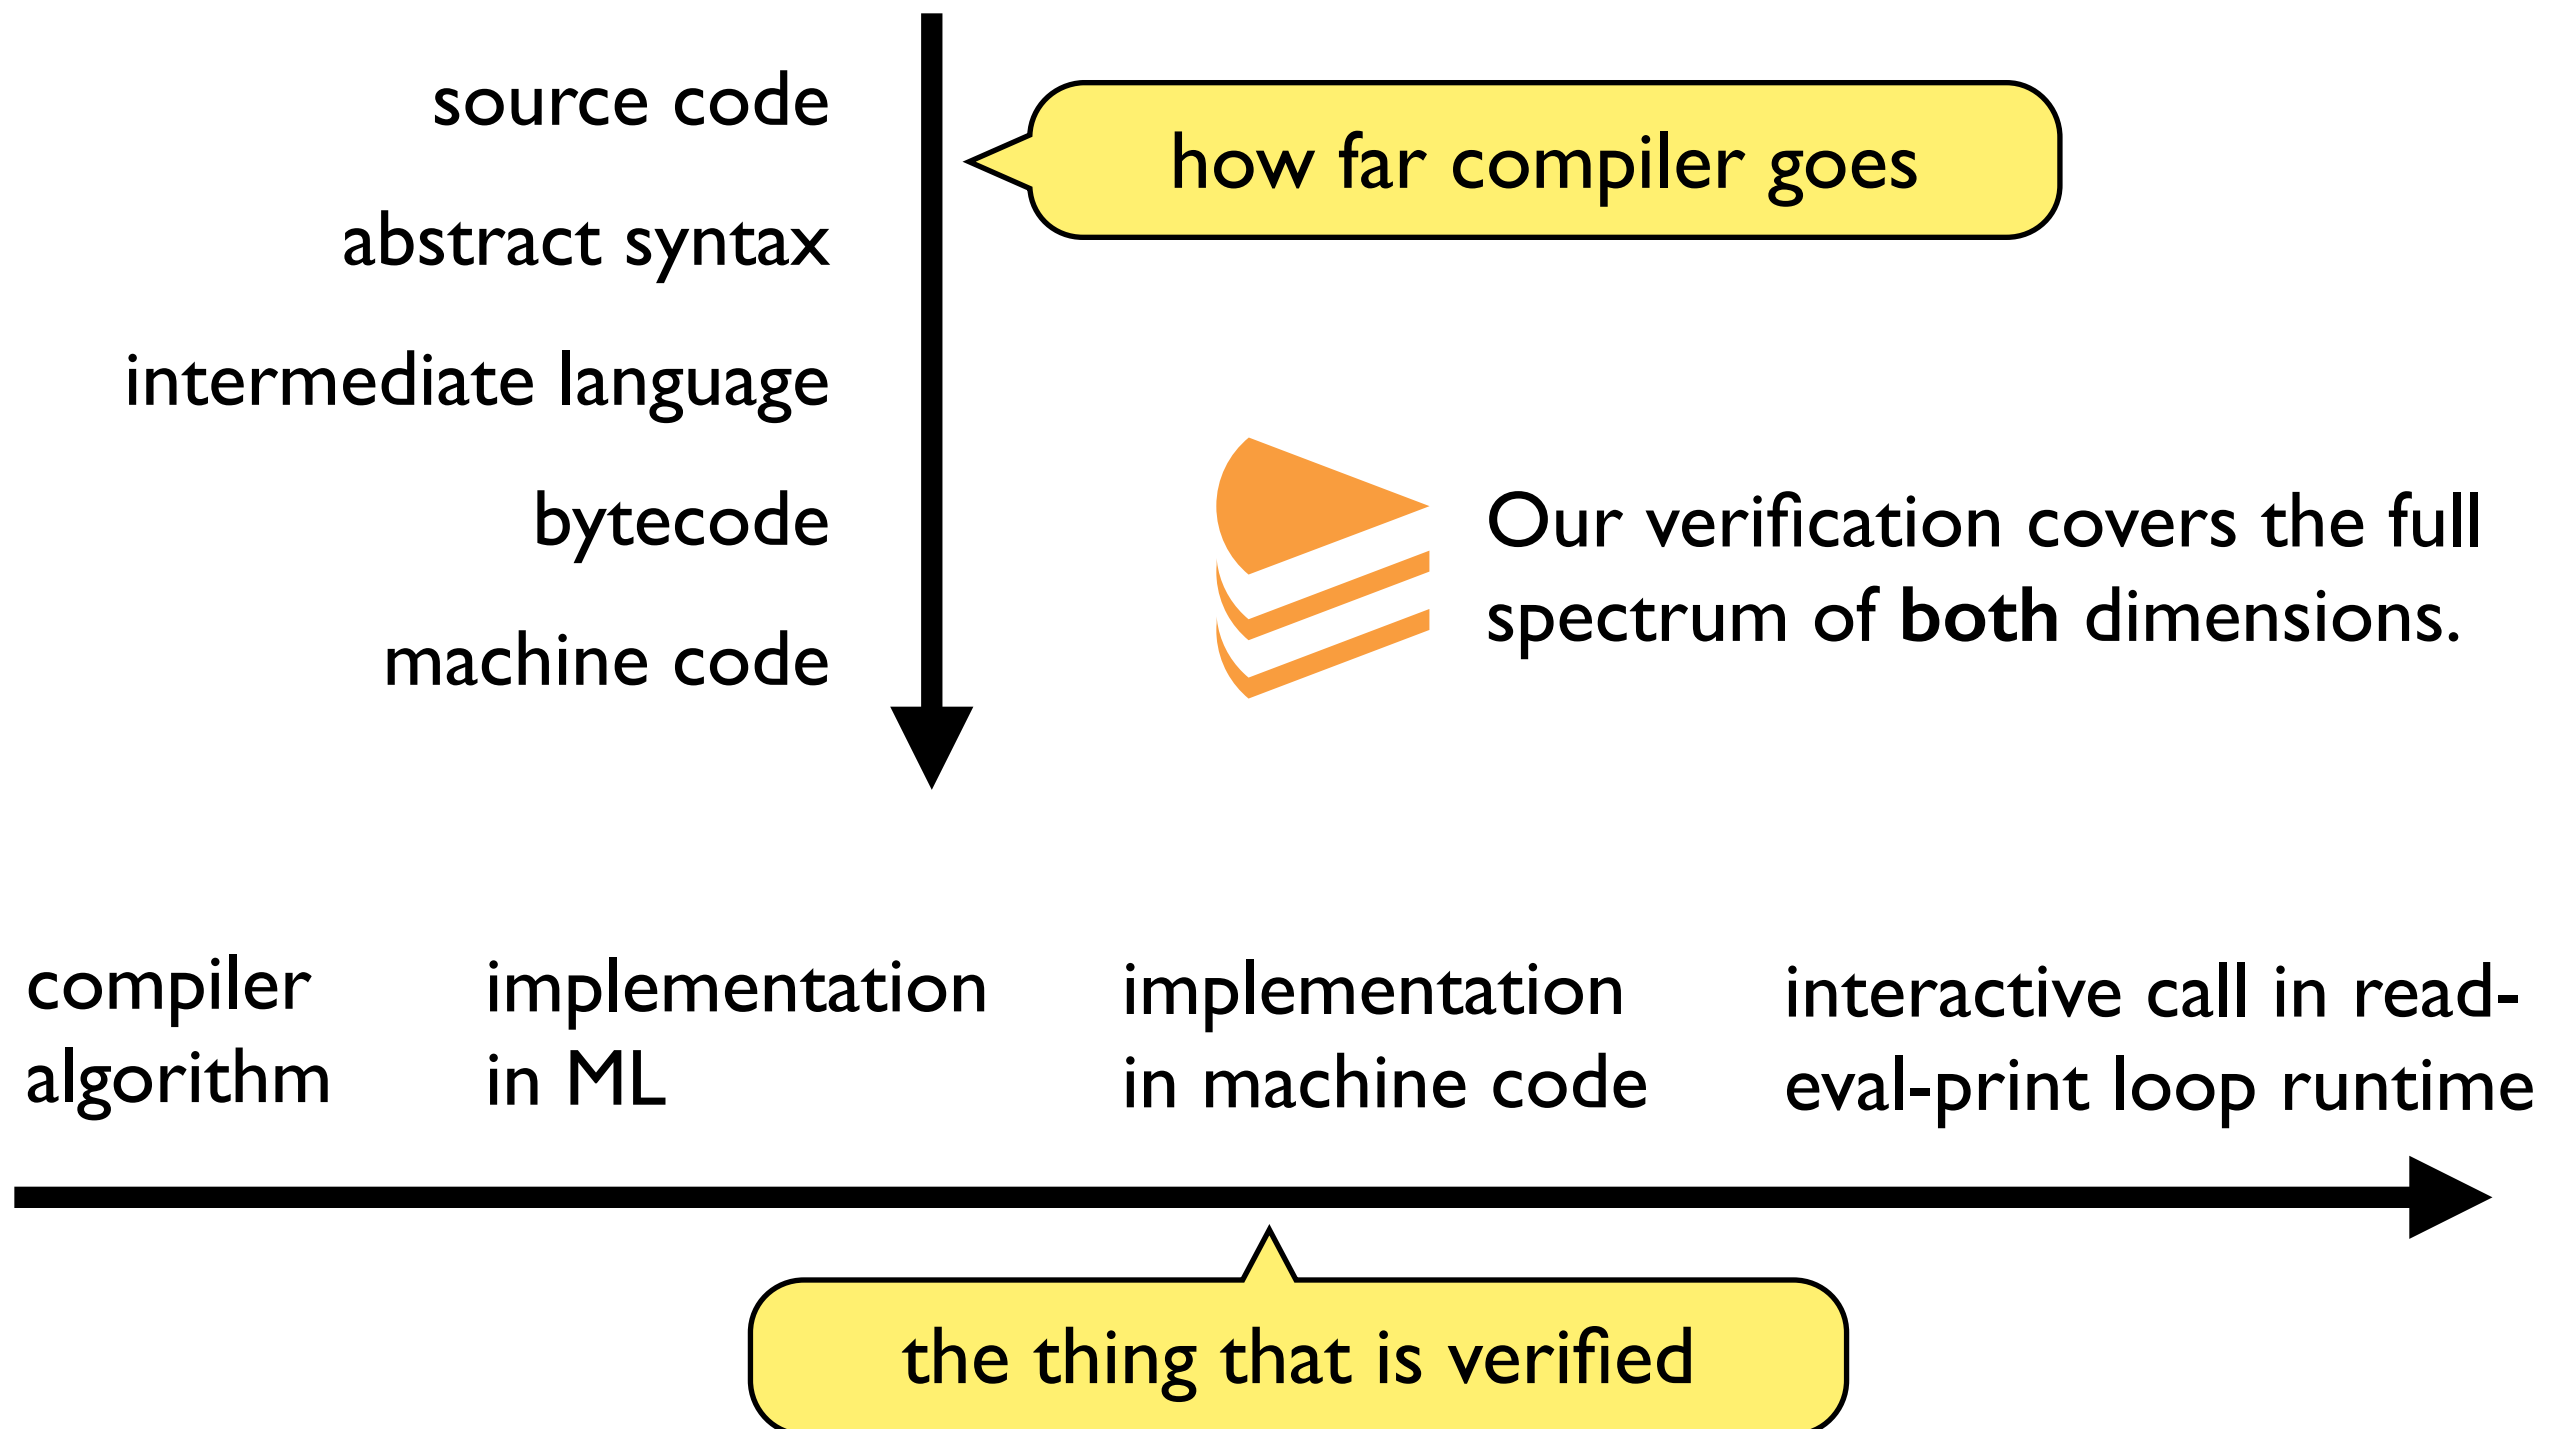
Abstract

We have developed and mechanically verified an ML system called CakeML, which supports a substantial subset of Standard ML. CakeML is implemented as an interactive read-eval-print loop that generates 32/64 machine code. Our correctness theorem ensures that only those results permitted by the semantics are produced on

1. Introduction

The last decade has seen a strong interest in verified compilation; and there have been significant, high-profile results, many based on the CompCert compiler for C [1, 14, 16, 29]. This interest is easy to justify: in the context of program verification, an unverified compiler forms a large and complex part of the trusted computing base. However, to our knowledge, none of the existing work on verified compilers for general-purpose languages has addressed all dimensions: one, the compilation of a list of

Dimensions of Compiler Verification



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
= Standard ML without I/O or 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

Contributions of POPL'14 paper

Artefacts

Specifications

Verified Algorithms

light-weight approach to
divergence preservation
with big-step op. sem.

Proof techniques

Divergence Preservation

Bootstrapping

use verified compiler to
produce **verified**
implementation of the compiler

Proof development where everything fits together.

Approach

Proof by refinement:

Step 1: specification of CakeML language

- ▶ big-step and small-step operational semantics

Step 2: functional implementation in logic

- ▶ read-eval-print-loop as verified function in logic

Step 3: production of verified x86-64 machine code

- ▶ produced mostly by bootstrapping the compiler

Operational semantics

Big-step semantics:

- ▶ big-step evaluation relation
- ▶ environment semantics (cf. substitution sem.)
- ▶ produces `TypeError` for badly typed evaluations (e.g. `1+nil`)
- ▶ stuck = divergence

Equivalent small-step semantics:

- ▶ used for type-soundness proof and definition of divergence

Read-eval-print-loop semantics.

Semantics written in Lem, see Mulligan et al. [ICFP'14]

Approach

Proof by refinement:

Step 1: specification of CakeML language

- ▶ big-step and small-step operational semantics

Step 2: functional implementation in logic

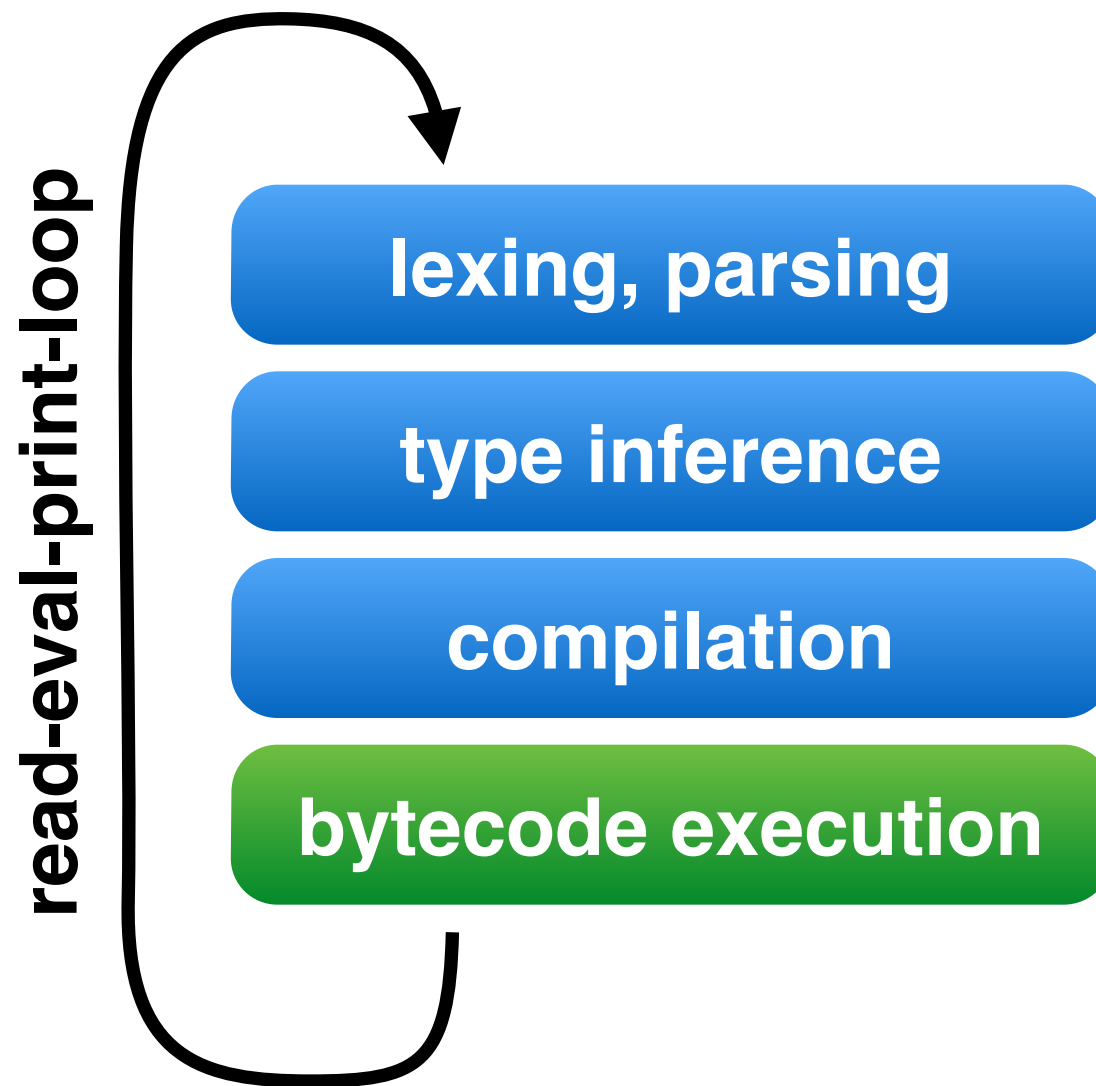
- ▶ read-eval-print-loop as verified function in logic

Step 3: production of verified x86-64 machine code

- ▶ produced mostly by bootstrapping the compiler

Functional implementation

Read-eval-print loop defined as rec. function in the logic:



lexing, parsing

Specification:

Context-free grammar (CFG) for significant subset of SML
Executable lexer.

Implementation:

Parsing-Expression-Grammar (PEG) Parser

- ▶ inductive evaluation relation
- ▶ executable interpreter for PEGs

Correctness:

Soundness and completeness

- ▶ induction on length of token list/parse tree and non-terminal rank

type inference

Specification:

Declarative type system.

Implementation:

Based on Milner's Algorithm W

Purely functional (uses state-exception monad)

Correctness:

new since last month!

Proved sound and *complete* w.r.t. declarative type system

Re-use of previous work on verified unification

compilation

Purpose:

Translates (typechecked) CakeML into CakeML Bytecode.

Implementation:

Translation via one intermediate language (IL).

- ▶ de Bruijn indices
- ▶ big-step operational semantics

CakeML to IL: makes language more uniform

IL to IL: removes pattern-matching, lightweight opt.

IL to Bytecode: closure conversion, data refinement, tail-call opt.

Semantics of **bytecode execution**

Instructions:

bc_inst	$::=$	Stack bc_stack_op PushExc PopExc
		Return CallPtr Call loc
		PushPtr loc Jump loc JumpIf loc
		Ref Deref Update Print PrintC char
		Label n Tick Stop
bc_stack_op	$::=$	Pop Pops n Shift $n\ n$ PushInt int
		Cons $n\ n$ El n TagEq n IsBlock n
		Load n Store n LoadRev n
		Equal Less Add Sub Mult Div Mod
loc	$::=$	Lab n Addr n

Small-step semantics; values and state:

bc_value	$::=$	Number int RefPtr n Block $n\ bc_value^*$
		CodePtr n StackPtr n
bc_state	$::= \{$	stack : bc_value^* ; refs : $n \mapsto bc_value$;
		code : bc_inst^* ; pc : n ; handler : n ;
		output : string; names : $n \mapsto$ string;
		clock : $n^?$ }

Semantics of **bytecode execution**

Sample rules:

$$\frac{\text{fetch}(bs) = \text{Stack } (\text{Cons } t \ n) \quad bs.\text{stack} = vs @ xs \quad |vs| = n}{bs \rightarrow (\text{bump } bs)\{\text{stack} = \text{Block } t \ (\text{rev } vs) :: xs\}}$$

$$\frac{\text{fetch}(bs) = \text{Return} \quad bs.\text{stack} = x :: \text{CodePtr } ptr :: xs}{bs \rightarrow bs\{\text{stack} = x :: xs; \text{pc} = ptr\}}$$

$$\frac{\text{fetch}(bs) = \text{CallPtr} \quad bs.\text{stack} = x :: \text{CodePtr } ptr :: xs}{bs \rightarrow bs\{\text{stack} = x :: \text{CodePtr } (\text{bump } bs).\text{pc} :: xs; \text{pc} = ptr\}}$$

compilation

Correctness:

Proved in the direction of compilation.

Shape of correctness theorem:

$$\begin{aligned} & (exp, env) \Downarrow_{ev} val \implies \\ & \text{“the code for } exp \text{ is installed in } bs \text{ etc.”} \implies \\ & \exists bs'. bs \rightarrow^* bs' \wedge \text{“} bs' \text{ contains } val \text{”} \end{aligned}$$

Bytecode semantics step relation

What about divergence?

We want: generated code diverges
if and only if source code diverges

Idea: add logical clock

Big-step semantics:

- has an optional **clock** component
- **clock 'ticks'** decrements every time a function is applied
- once **clock hits zero**, execution stops with a **TimeOut**

Why do this?

- because now big-step semantics describes both terminating and non-terminating evaluations

for every exp env $clock$

there is some result

either: Result
or TimeOut

$\forall exp\ env\ clock. \exists res. (exp, env, \text{Some } clock) \Downarrow_{ev} res$

produced by the semantics

Divergence

Evaluation diverges if

$$\forall clock. (exp, env, \text{Some } clock) \Downarrow_{\text{ev}} \text{TimeOut}$$

for all clock values

TimeOut happens

Compiler correctness proved in conventional forward direction:

$$(exp, env) \Downarrow_{\text{ev}} val \implies$$

“the code for *exp* is installed in *bs* etc.” \implies

$$\exists bs'. bs \rightarrow^* bs' \wedge \text{“} bs' \text{ contains } val \text{”}$$

Bytecode has clock

... that stays in sync with CakeML clock

Theorem: *bytecode diverges* if and only if *CakeML eval diverges*

Approach

Proof by refinement:

Step 1: specification of CakeML language

- ▶ big-step and small-step operational semantics

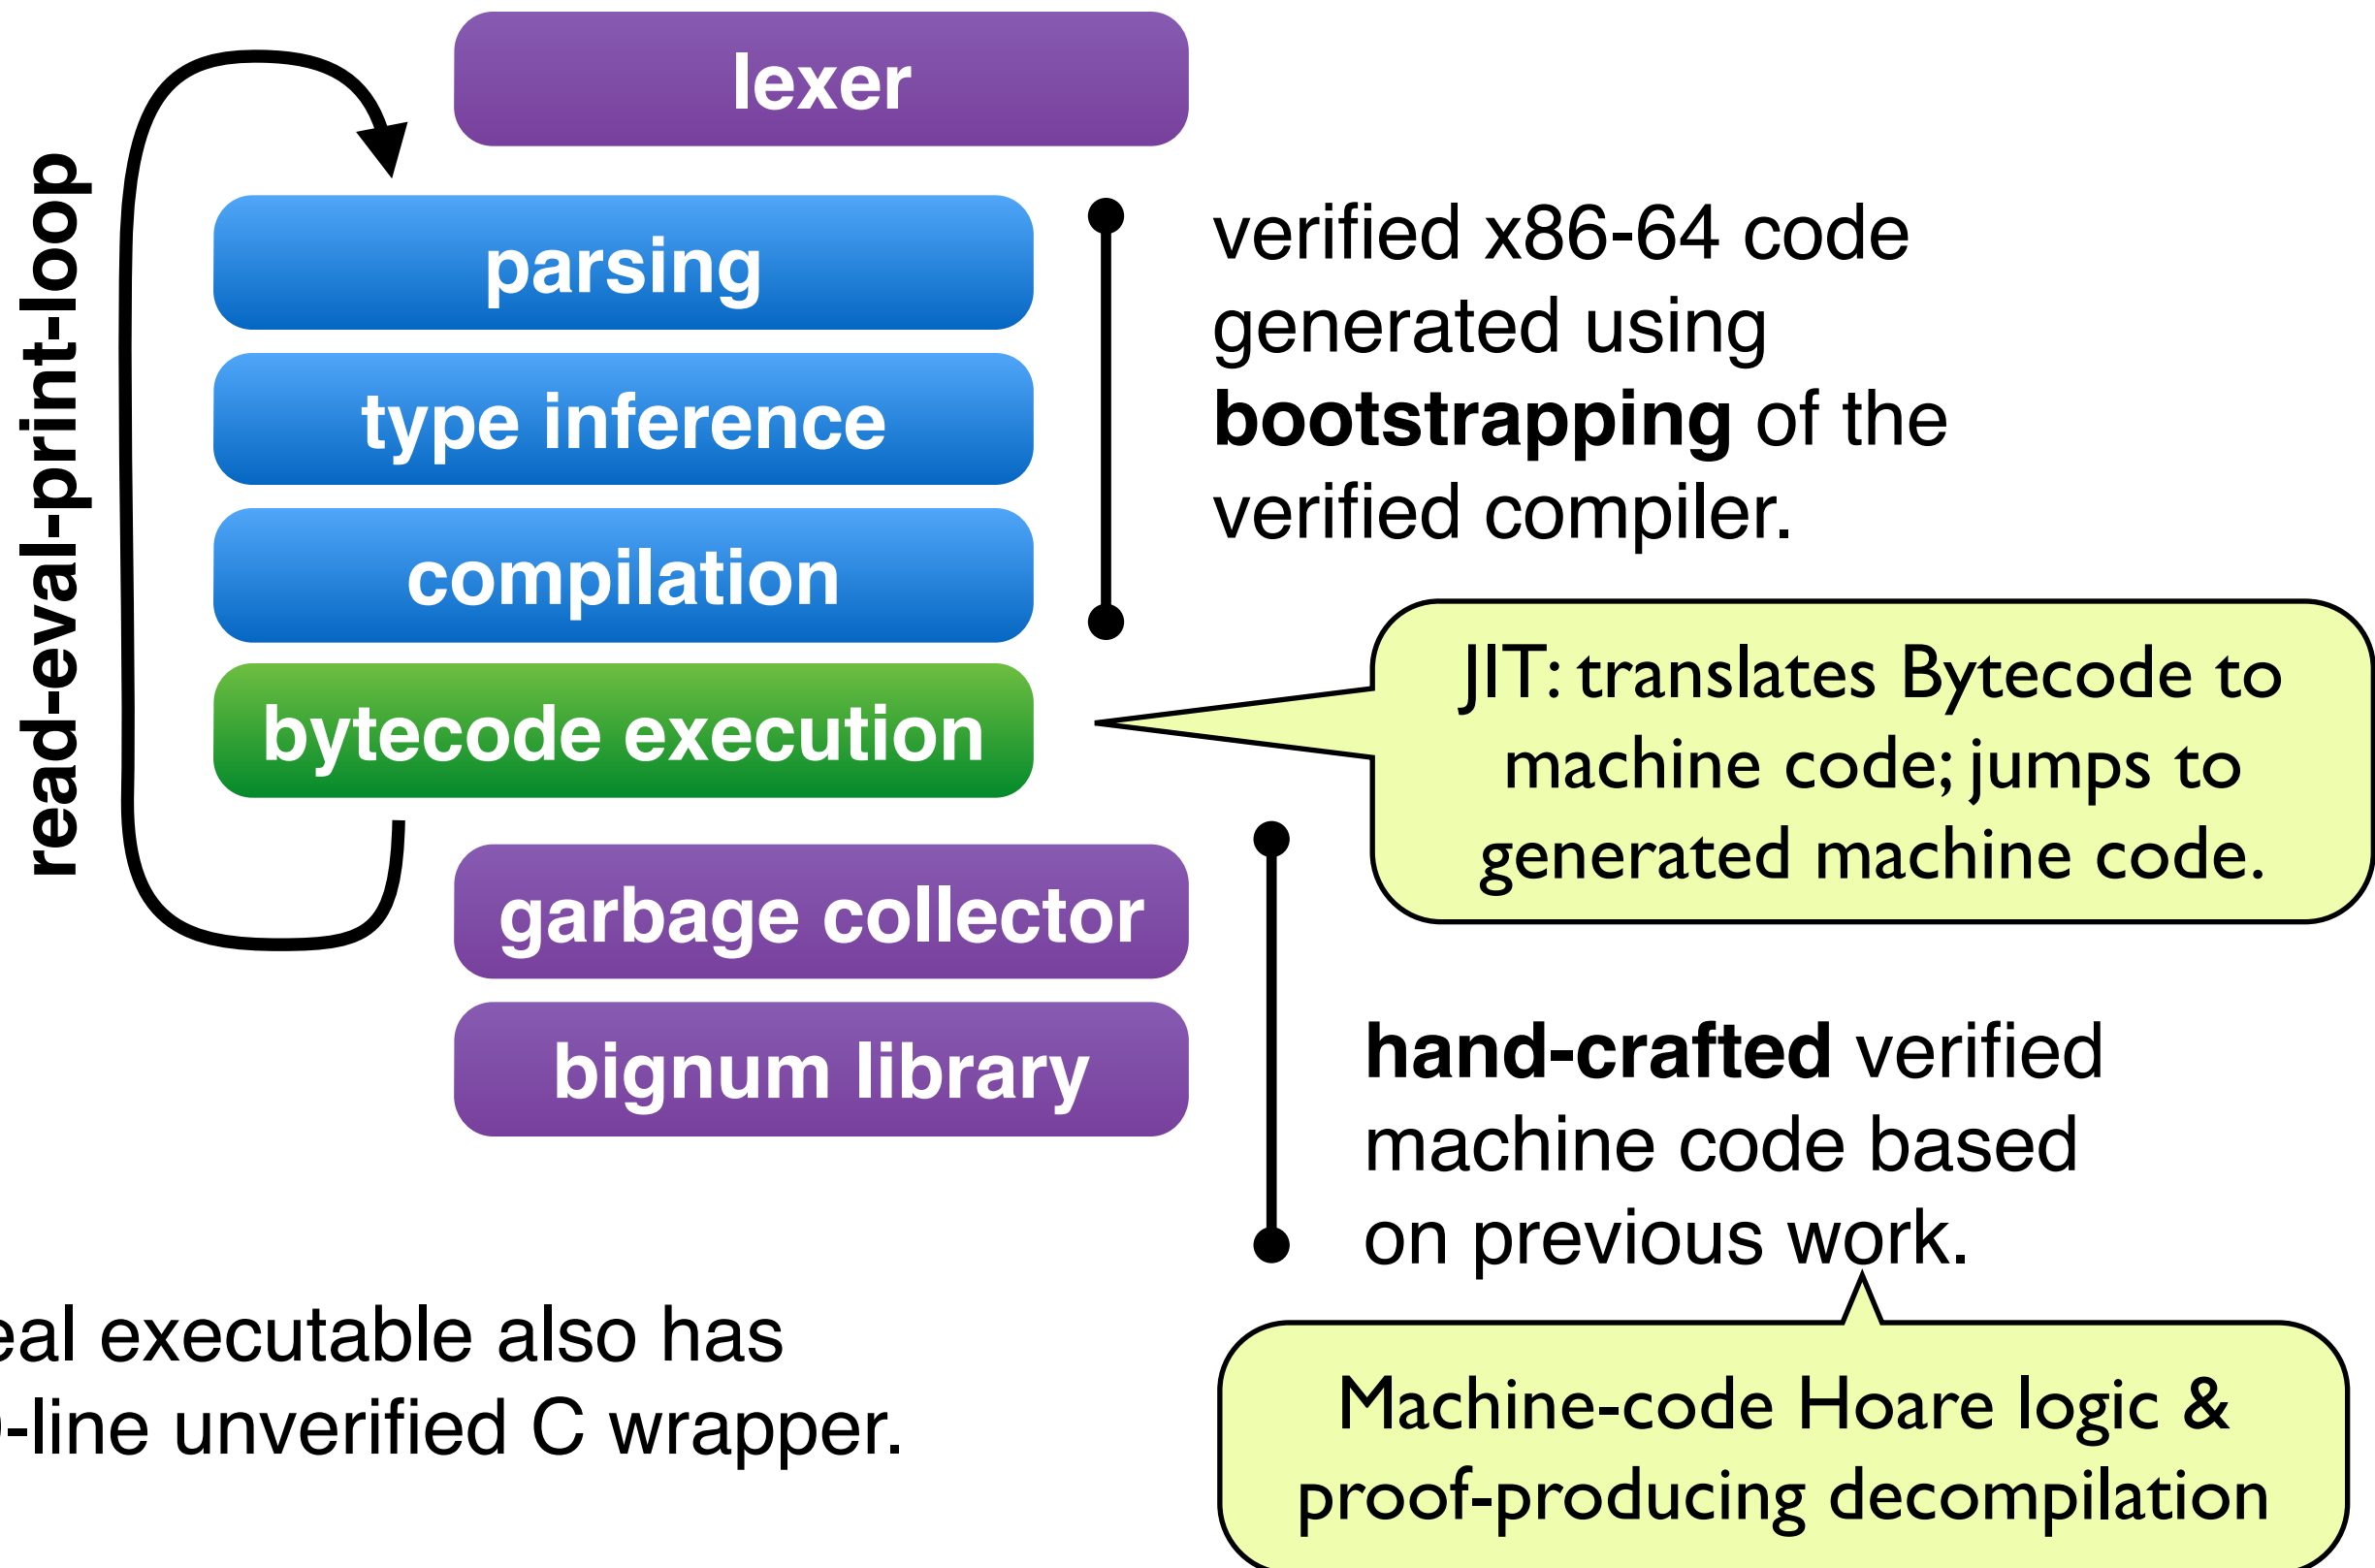
Step 2: functional implementation in logic

- ▶ read-eval-print-loop as verified function in logic

Step 3: production of **verified x86-64 machine code**

- ▶ produced mostly by bootstrapping the compiler

Verified x86-64 machine code



Compiling Bytecode into x86-64 mc

Extract of definition:

each bytecode inst. maps to some x86

$\text{x64 i Pop} = [0\text{x}48, 0\text{x}58]$

$\text{x64_code i (x::xs)} = \text{x64 i x} ++ \text{x64_code (i + len(x64 i x)) xs}$

Correctness:

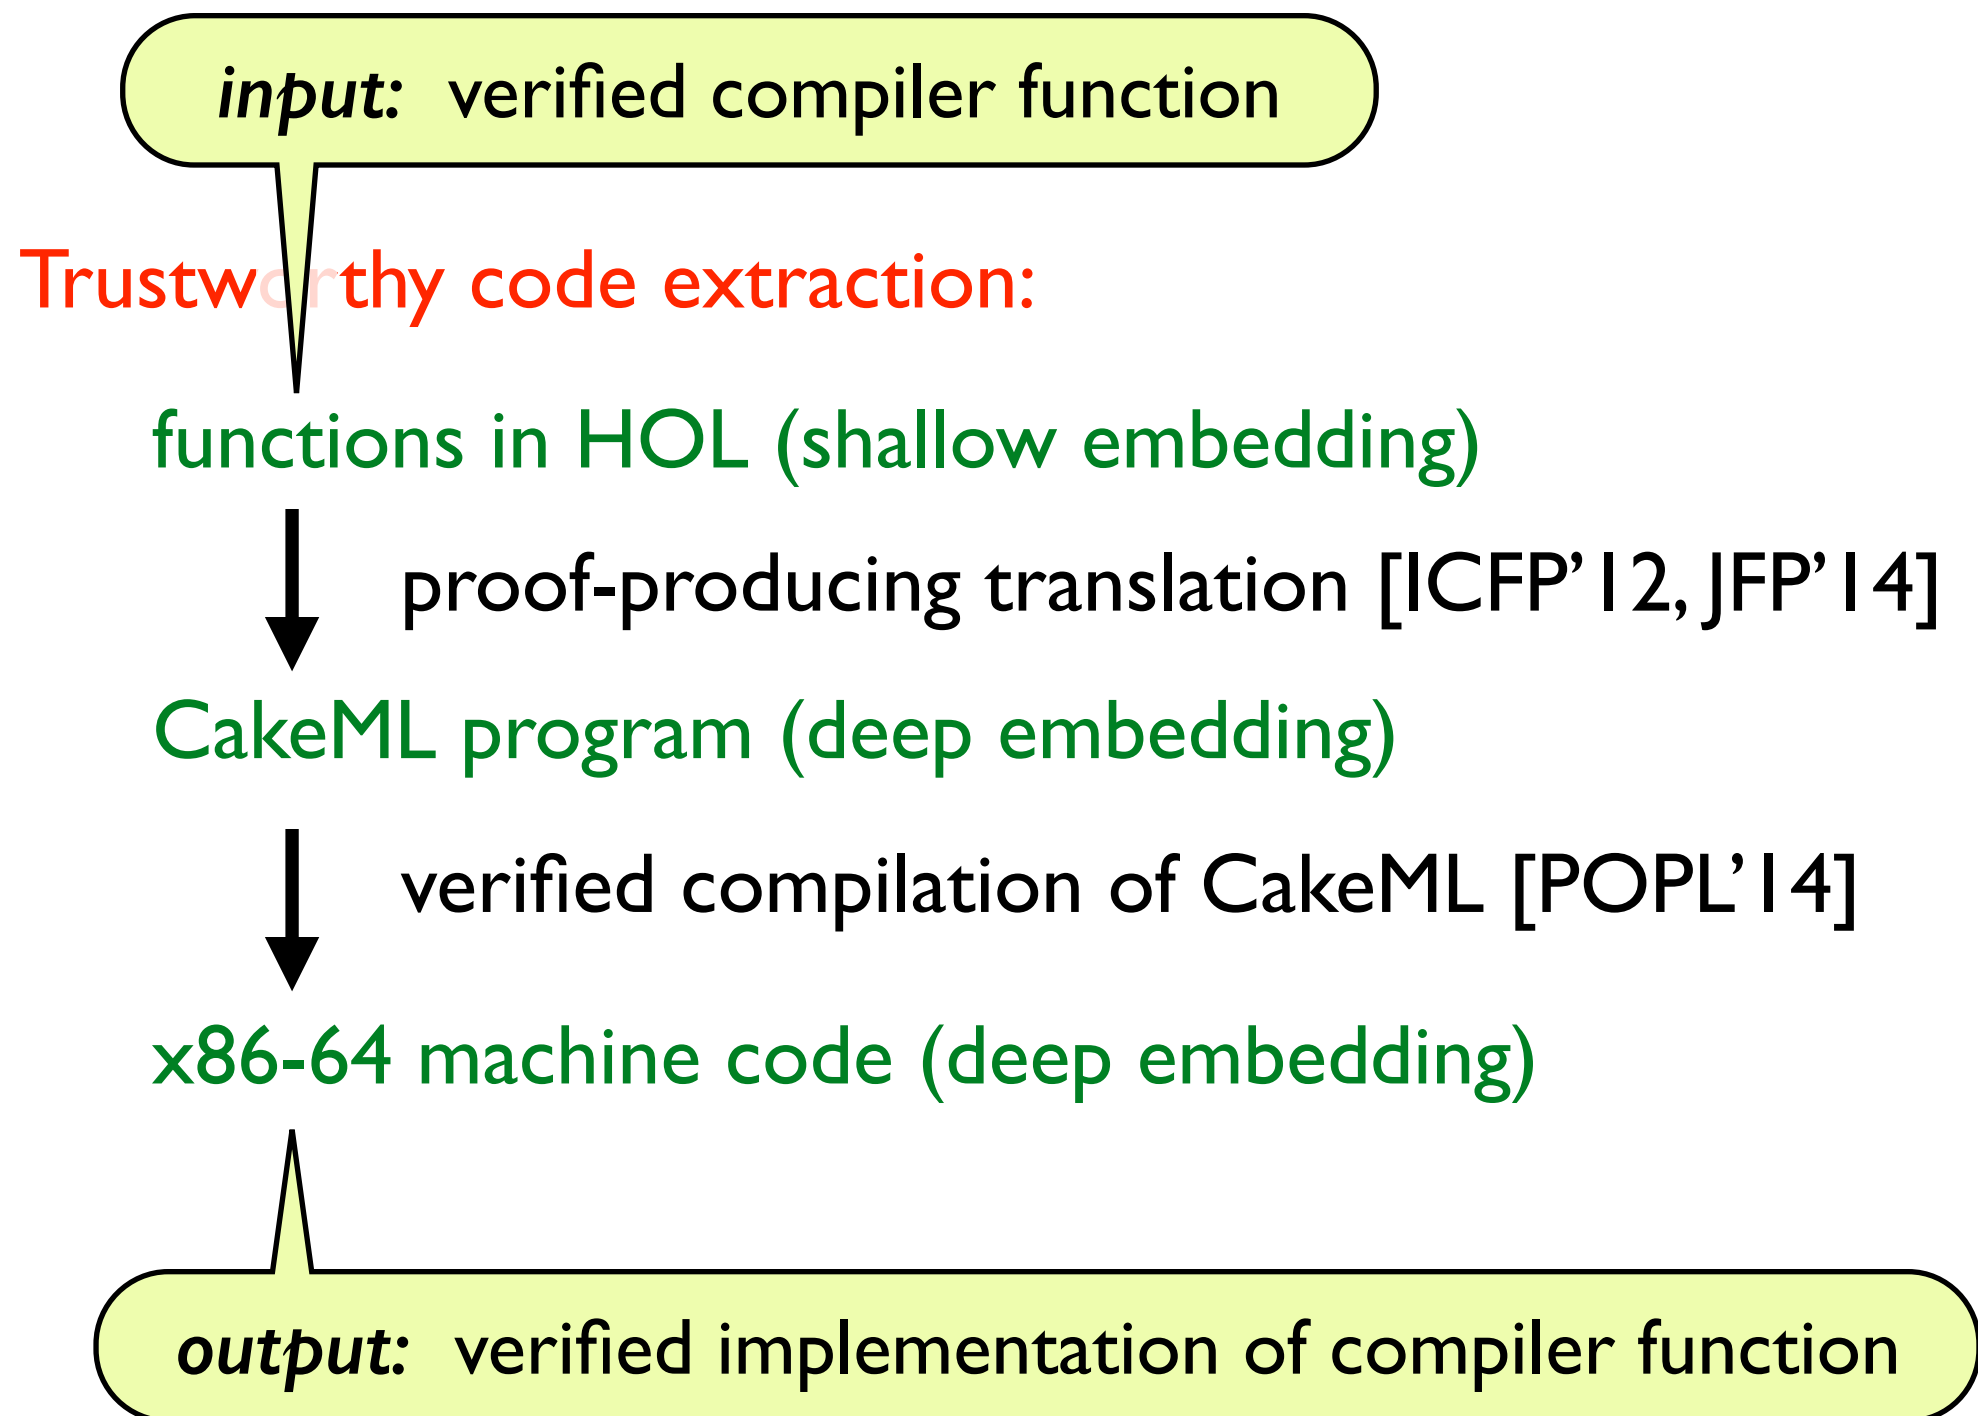
Each Bytecode instruction is correctly executed by generated x86-64 code.

$bs \rightarrow bs' \implies$
 $\text{temporal } \{(base, \text{x64_code } 0 \text{ } bs.\text{code})\}$
 $(\text{now } (bc_heap \text{ } bs \text{ } (base, aux)) \implies$
 $\quad \text{later } (\text{now } (bc_heap \text{ } bs' \text{ } (base, aux))$
 $\quad \vee \text{ now } (\text{out_of_memory_error } aux)))$


heap invariant / memory abstraction

Bootstrapping the verified compiler

Idea for in-logic bootstrapping



Compiling the compiler in logic



function in logic: **compile**



by proof-producing synthesis [ICFP'12]

CakeML program (**COMPILE**) such that:
 \vdash **COMPILE** *implements* compile

Proof by evaluation inside the logic:

\vdash compile-to-x64 **COMPILE** = x64-code

Compiler correctness theorem:

$\vdash \forall \text{prog. compile-to-x64 prog}$ *implements* prog

Combination of theorems:

\vdash x64-code *implements* compile

Details (build up)

Function in logic:

$$\begin{aligned}\text{map } f \ [] &= [] \\ \text{map } f \ (h::t) &= f \ h::\text{map } f \ t\end{aligned}$$

Translation into CakeML produces:

```
fun map v3 v4 =  
  case v4  
  of [] => []  
   | v2::v1 => v3 v2::(map v3 v1)
```

Evaluation in logic:

$$\vdash \text{map length } [[1; 1]; [2]; []] = [2; 1; 0]$$

Translation into CakeML, actual output:

```
map_dec =  
  Letrec  
    [ ("map", "v3",  
      Fun "v4"  
        (Mat (Var "v4")  
          [(Pcon "nil" [], Con "nil" []);  
           (Pcon "::" [Pvar "v2"; Pvar "v1"],  
            Con "::"  
              [App [Var "v3"; Var "v2"]; App [App [Var "map"; Var "v3"]; Var "v1"]]))]))]
```

Details (build up)

Produced proof (called a certificate theorem):

$\vdash \exists env\ c.$

$\text{EvalDec InitEnv map_dec } env \wedge \text{Lookup "map" } env = \text{Some } c \wedge$
 $((a \longrightarrow b) \longrightarrow \text{ListTy } a \longrightarrow \text{ListTy } b) \text{ map } c$

Translation into CakeML, actual output:

```
map_dec =  
  Letrec  
    [ ("map", "v3",  
      Fun "v4"  
        (Mat (Var "v4")  
          [(Pcon "nil" [], Con "nil" []);  
           (Pcon "::" [Pvar "v2"; Pvar "v1"],  
            Con "::"  
              [App [Var "v3"; Var "v2"]; App [App [Var "map"; Var "v3"]; Var "v1"]])])]) ]
```


Evaluation of compilation (in logic)

Example 2 (Compilation by evaluation of map)

$\vdash \text{CompileDec InitCS map_dec} =$
 $(\text{MapCS},$
 $[\text{Jump (Lab 12); Label 10; Stack (PushInt 0); Stack (PushInt 1); Ref};$
 $\text{PushPtr (Lab 11); Stack (Load 0); Stack (Load 5); Stack (PushInt 1);}$
 $\text{Stack (Cons 0); Stack (... ..); ; ... ; ... }])$

Compiler correctness (specialised to terminating case):

$\vdash \text{Inv env}_1 \text{ cs}_1 \text{ bs}_1 \wedge \text{EvalDec env}_1 \text{ dec env}_2 \wedge \text{CompileDec cs}_1 \text{ dec} = (\text{cs}_2, \text{bc}) \Rightarrow$
 $\exists \text{bs}_2. (\text{AddCode bs}_1 \text{ bc}) \rightarrow^* \text{bs}_2 \wedge \text{Halted bs}_2 \wedge \text{Inv env}_2 \text{ cs}_2 \text{ bs}_2$

(Some of the) actual details

Translation of compiler into CakeML

$\vdash \exists c.$
 $\text{EvalDec InitEnv (Struct "C" CompileDec_decs) CompEnv} \wedge$
 $\text{LookupMod "C" "compiledec" CompEnv} = \text{Some } c \wedge$
 $(\text{CompStateTy} \longrightarrow \text{DecTy} \longrightarrow \text{PairTy CompStateTy (ListTy BCInstTy)}) \text{ CompileDec } c$

Evaluating the compiler on itself

$\vdash \text{CompileDec InitCS (Struct "C" CompileDec_decs)} =$
 $(\text{CompCS}, \text{CompileDec_bytecode})$

Compiler correctness theorem:

$\vdash \text{Inv InitEnv InitCS InitBS}$
 $\vdash \text{Inv } env_1 \text{ } cs_1 \text{ } bs_1 \wedge \text{EvalDec } env_1 \text{ } dec \text{ } env_2 \wedge \text{CompileDec } cs_1 \text{ } dec = (cs_2, bc) \Rightarrow$
 $\exists bs_2. (\text{AddCode } bs_1 \text{ } bc) \rightarrow^* bs_2 \wedge \text{Halted } bs_2 \wedge \text{Inv } env_2 \text{ } cs_2 \text{ } bs_2$

NB: For a read-eval-print-loop, the details are a bit more involved...

Top-level **correctness theorem**

Top-level correctness theorem

Top-level specification:

initial program, the basis library

ReplSem Basis *input output*

string

string ending in either
Terminate or Diverge

Correctness theorem:

$\vdash \text{TemporalX64 ReplX64}$

$(\text{Now} (\text{InitialisedX64 } ms) \Rightarrow$

$\diamond \text{Now} (\text{OutOfMemX64 } ms) \vee$

$\exists \text{output}.$

$\text{Holds} (\text{ReplSem Basis } ms.\text{input } output) \wedge$

$\text{if Diverges } output \text{ then } \Box \diamond \text{Now} (\text{RunningX64 } output \ ms)$

$\text{else } \diamond \text{Now} (\text{TerminatedX64 } output \ ms))$

Numbers

Performance:

Slow: *interpreted* OCaml is 1x faster (... future work!)

Effort:

~100k lines of proof script in HOL4

< 5 man-years, but builds on a lot of previous work

Size:

875,812 instructions of verified x86-64 machine code

implementation generates
more instructions at runtime

large due to bootstrapping, naive compiler

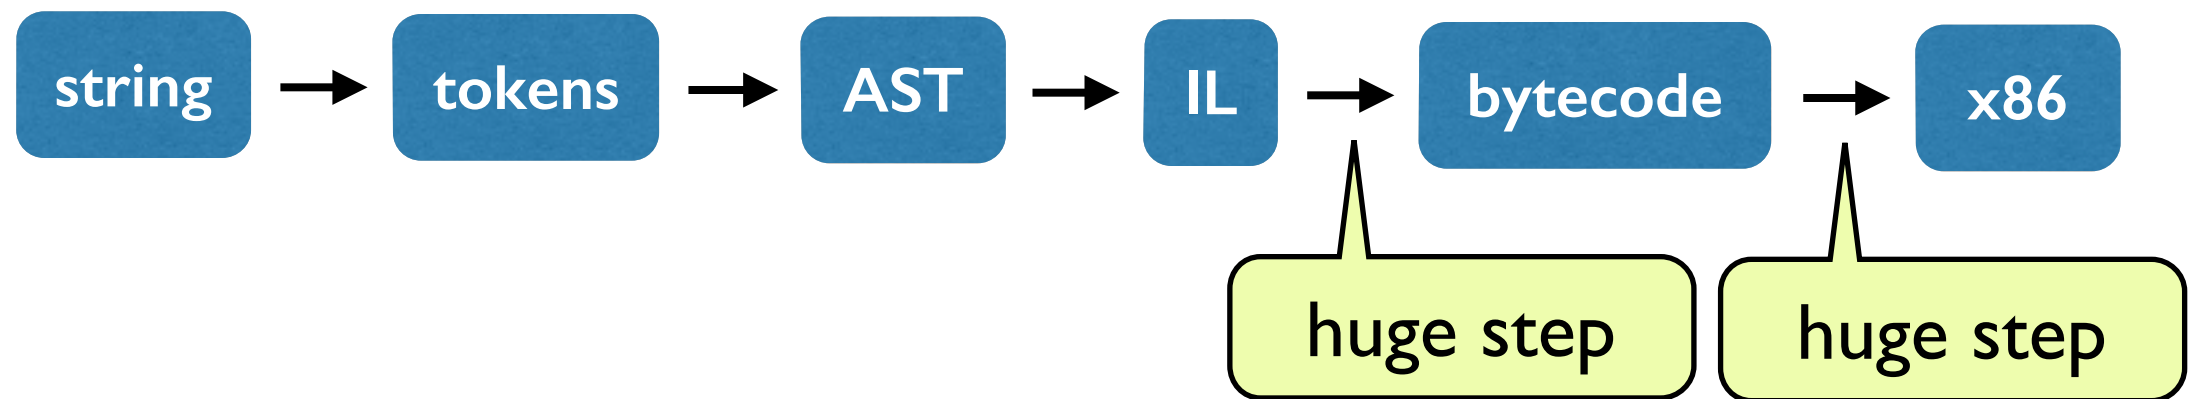
This talk

Part 1: verified implementation of CakeML

Part 2: current status, HOL light, future

Current status

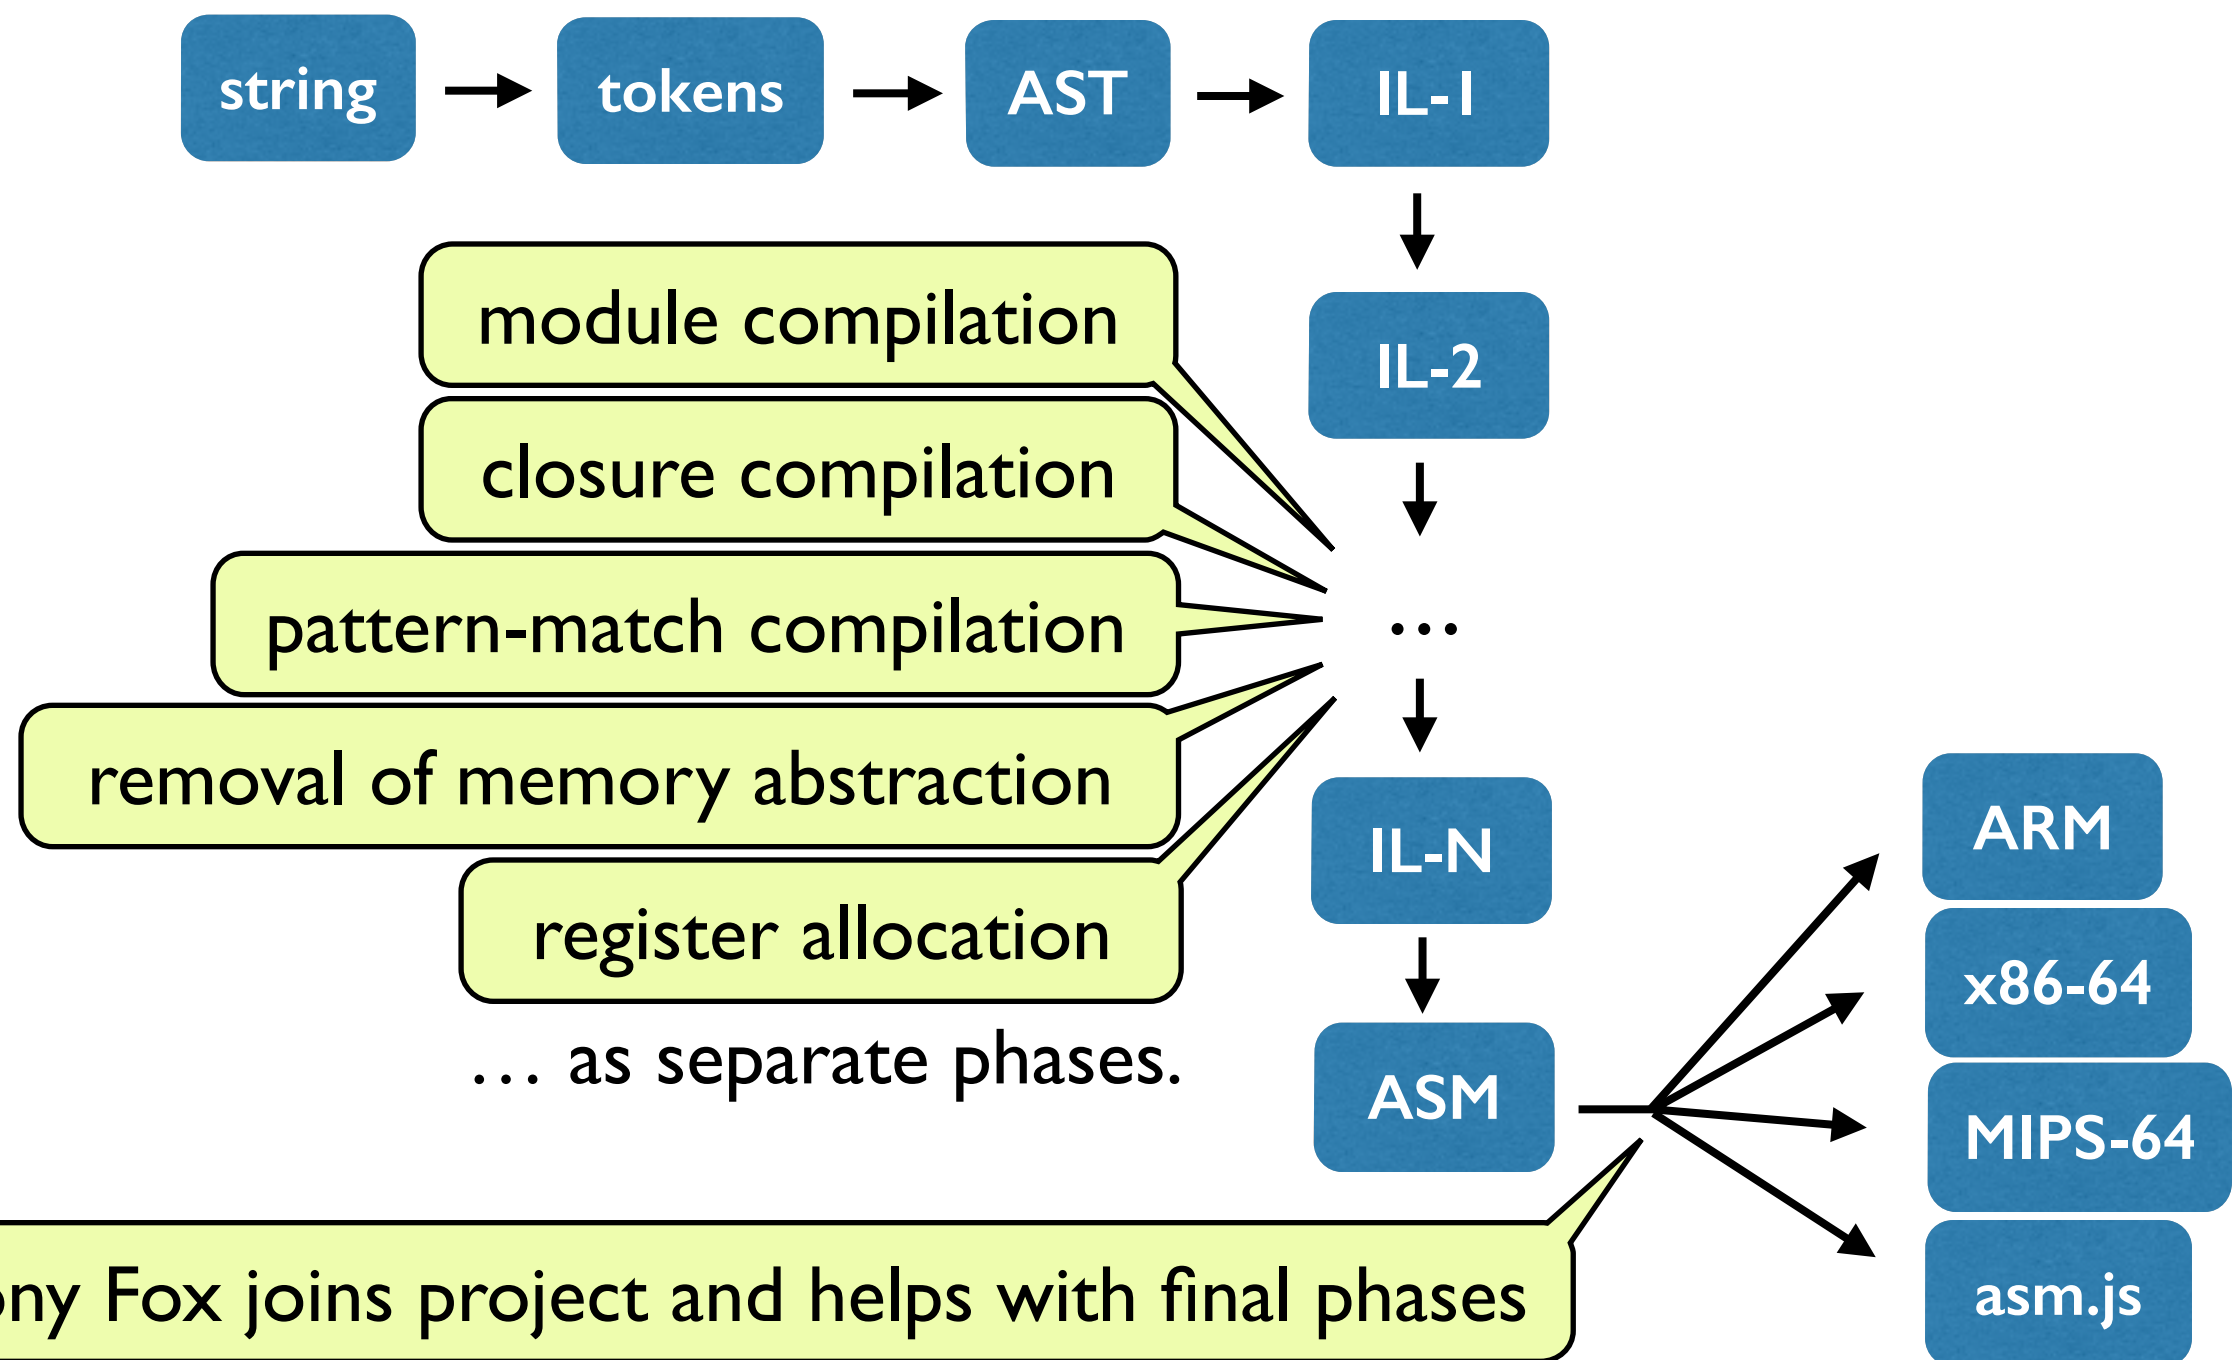
Current compiler:



Bytecode simplified proofs of read-eval-print loop, but made optimisation impossible.

Current work / future plans

Refactored compiler: split into more conventional compiler phases



Verified examples on CakeML

Verification infrastructure:

- have: synthesis tool that maps HOL into CakeML [ICFP'12, JFP'14]
- future: integration with Arthur Charguéraud's characteristic formulae technology [ICFP'10, ICFP'11]

for developing interesting verified examples.

Big example: verified HOL light

ML was originally developed to host theorem provers.

Aim: verified HOL theorem prover.

We have [ITP'13, ITP'14]:

- syntax, semantics and soundness of HOL (stateful, stateless)
- verified implementation of the HOL light kernel in CakeML (produced through synthesis)

Still to do:

- soundness of kernel \Rightarrow soundness of entire HOL light
- run HOL light standard library on top of CakeML

Freek Wiedijk is translating HOL light sources to CakeML

Summary

Contributions so far:

First **bootstrapping** of a formally **verified compiler**.

New lightweight method for **divergence preservation**.

Current work:

Formally **verified** implementation of **HOL light**.

Verified **I/O** (foreign-function interface). **seL4**.

Compiler improvements (new ILs, opt, targets).

Long-term aim:

An **ecosystem** of tools and proofs around CakeML lang.

Questions? **Suggestions?**