



CakeML

From functions to machine code
... with proof all the way!

Speaker: Magnus Myreen

DeepSpec workshop — PLDI 2018

The CakeML project

Main goals:

An *ITP as a programming environment* with proofs.

Provide a *code extraction mechanism* that relates *via proof* the extracted (machine) code with the original HOL functions.

By-product:

A *realistic verified ML compiler* built as a base for research and teaching.

Example

program that computes frequencies of words

```
hello! said the sailor  
the captain and the  
navigator replied hello!
```



```
and: 1  
captain: 1  
hello!: 2  
navigator: 1  
replied: 1  
said: 1  
sailor: 1  
the: 3
```

Specification

There is some list of words (ws) that matches the words in the file ...

... the list is sorted ...

$\text{valid_wordfreq_output } file_contents \text{ output} \iff$

$\exists ws.$

$\text{set } ws = \text{set } (\text{words_of } file_contents) \wedge \text{sorted } (\lambda x y. x < y) ws \wedge$
 $\text{output} = \text{concat } (\text{map } (\lambda w. \text{format_output } (w, \text{frequency } file_contents w)) ws)$

... and the output has a line per word indicating the frequency of the word.

Pure implementation

Simple functional program in the logic of an ITP:

```
compute_wordfreq_output input_lines =  
map format_output (to_list (foldl insert_line empty_tree input_lines))  
insert_line t line = foldl insert_word t (words_of line)
```

idea: builds an ordered binary tree and then flattens it

An easy-to-prove correctness theorem:

```
⊢ valid_wordfreq_output file_contents  
  (concat (compute_wordfreq_output (lines_of file_contents)))
```

Monads for I/O part

```
wordfreq () =  
do  
  args ← commandline (arguments ());  
  filename ← hd args;  
  lines ← stdio (inputLinesFrom filename);  
  stdio (print_list (compute_wordfreq_output lines))  
od otherwise  
do  
  name ← commandline (name ());  
  stdio (print_err (“usage: ” ^ name ^ “ filename\n”))  
od
```

monadic function for I/O in “Haskell-style”

Verification easy thanks to straightforward shallow embedding.

Correctness of compiler output

most compiler verification projects promise this, but few deliver fully

The CakeML toolchain allows *transportation* of correctness properties to the compiler output.

CakeML delivers:

theorem proved inside an ITP

$$\begin{aligned} &\vdash \text{wfCL } [pname; fname] \wedge \text{wfFS } fs \wedge \text{hasFreeFD } fs \wedge \\ &\text{get_file_contents } fs \text{ } fname = \text{Some } file_contents \wedge \\ &x64_installed \text{ compiler_output } (\text{basis_ffi } [pname; fname] fs) \text{ } mc \text{ } ms \Rightarrow \\ &\exists io_events \text{ } ascii_output. \\ &\text{machine_sem } mc (\text{basis_ffi } [pname; fname] fs) \text{ } ms \subseteq \\ &\text{extend_with_resource_limit } \{ \text{Terminate Success } io_events \} \wedge \\ &\text{extract_fs } fs \text{ } io_events = \text{Some } (\text{add_stdout } fs \text{ } ascii_output) \wedge \\ &\text{valid_wordfreq_output } file_contents \text{ } ascii_output \end{aligned}$$

Toolchain

for proof-producing code extraction

Components of toolchain

function in the logic



AST



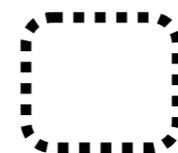
compiler backend



machine code

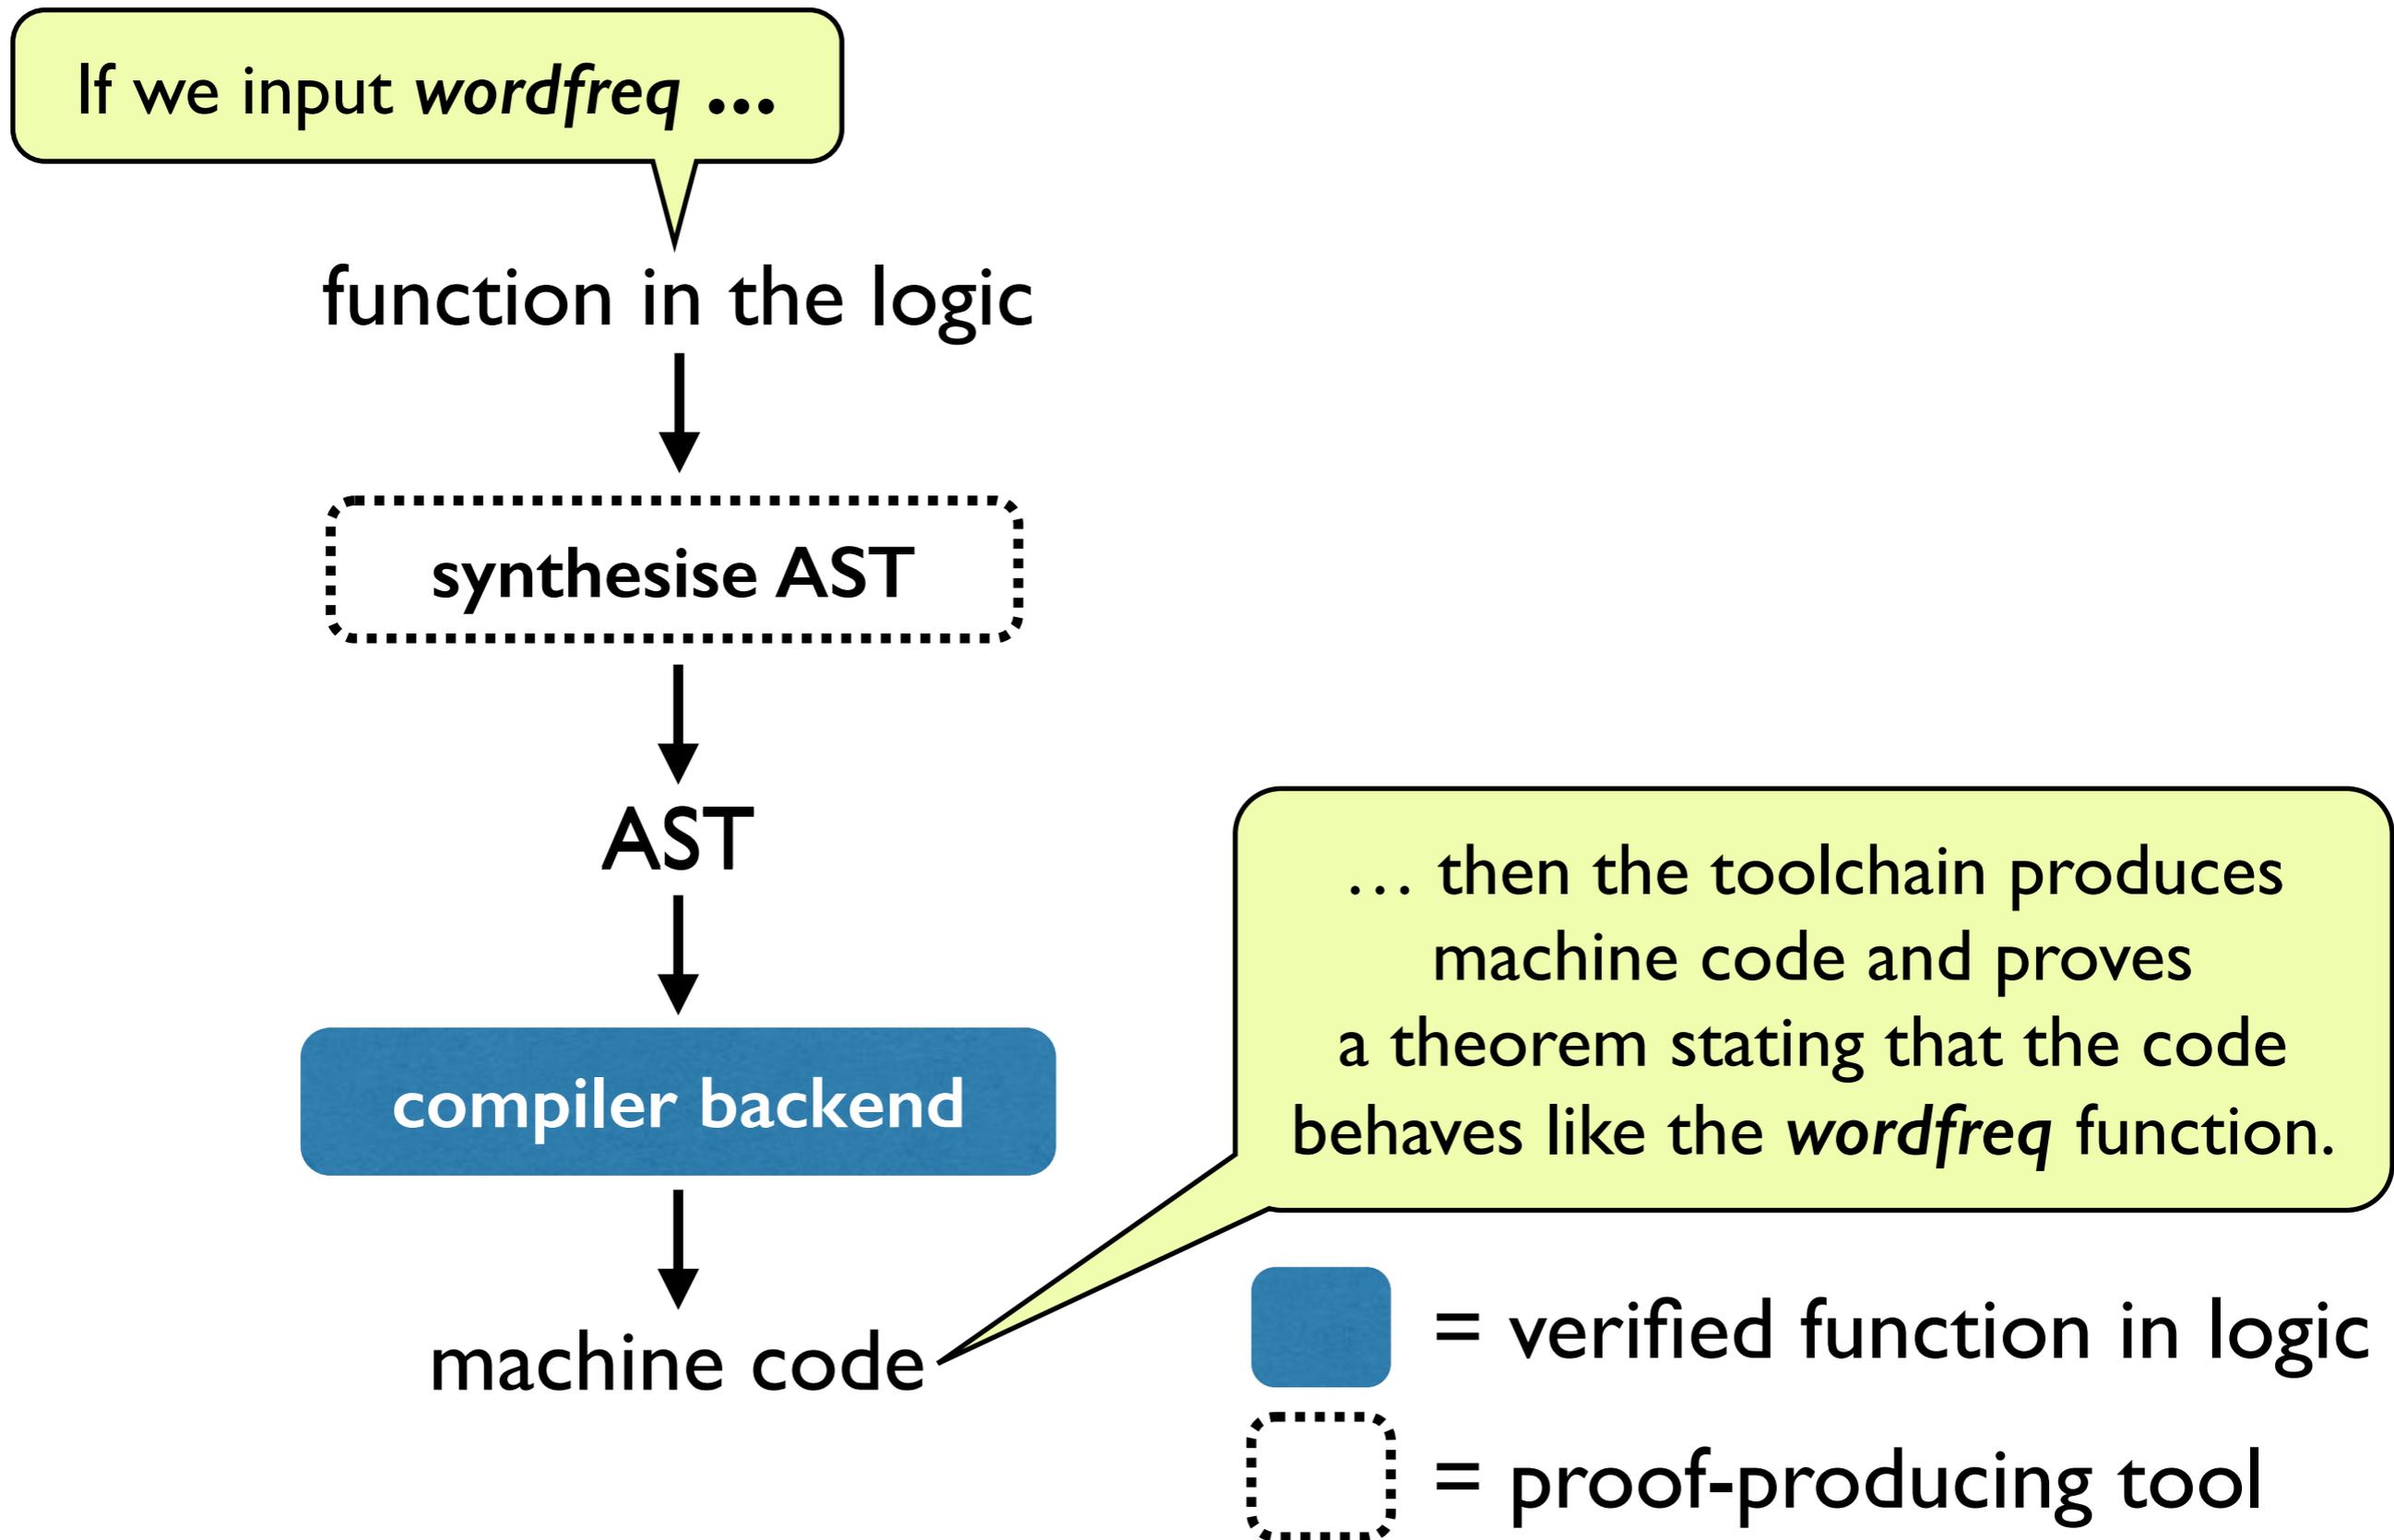


= verified function in logic



= proof-producing tool

Components of toolchain



a function in the logic

An aside:

Compiler Bootstrapping

concrete syntax



SML parser



type inferencer



AST



compiler backend

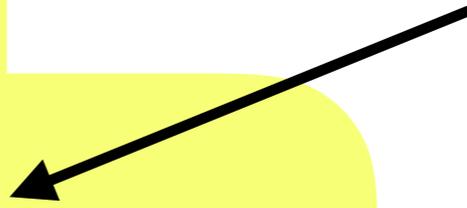


machine code

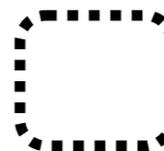
function in the logic



synthesise AST



= verified function in logic



= proof-producing tool

An aside: *Compiler Bootstrapping*

If we input
the CakeML compiler ...

function in the logic

synthesise AST

AST

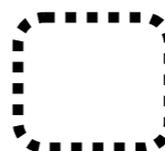
compiler backend

machine code

... then gen. code behaves like
the CakeML compiler function.

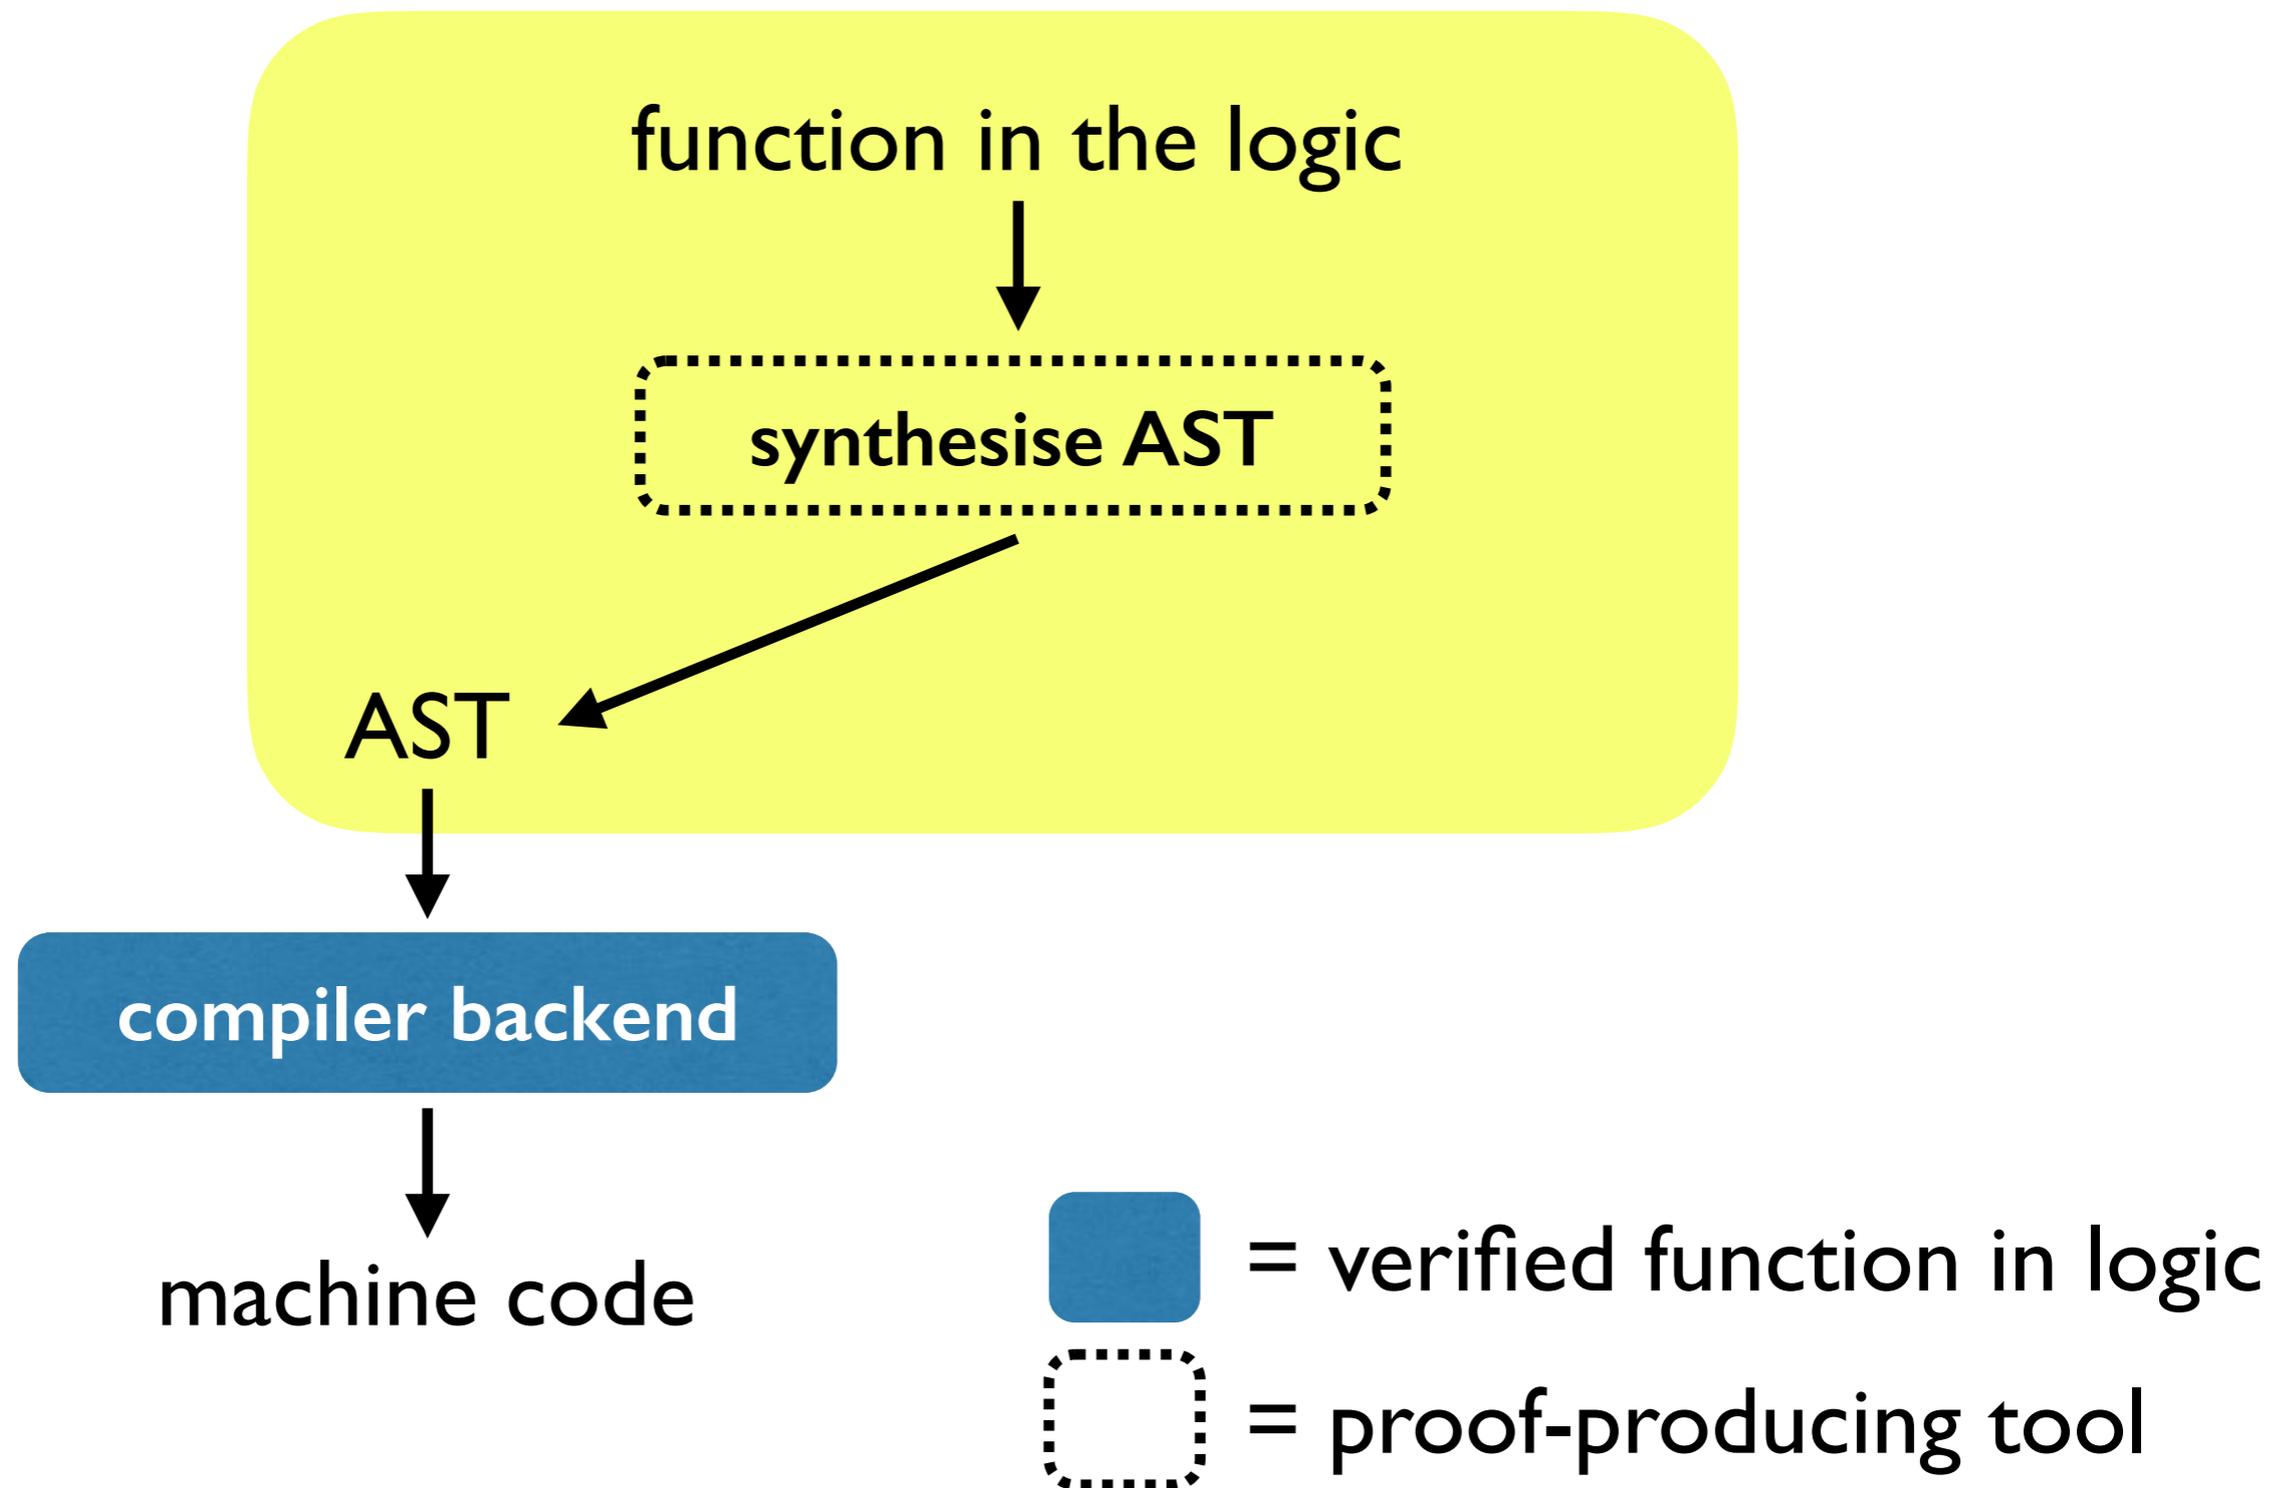


= verified function in logic



= proof-producing tool

Proof-producing synthesis



Proof-producing synthesis

Key definition:

CakeML exp evaluates to value v

$$\text{Eval } env \ exp \ post = \exists v. \langle env, exp \rangle \Downarrow v \wedge post \ v$$

big-step semantics, total-correctness

Example:

This states: CakeML expression 1 relates to ...

$$\text{Eval } env \ [1] \ (\text{int } 1) \ \dots \ \text{integer } 1 \ \text{in the logic.}$$

where $\text{int } i = (\lambda v. v = \text{Litv } (\text{IntLit } i))$

Automation

Each stage automation proves:

$assumptions \Rightarrow Eval\ env\ code\ (ref_inv\ t)$

Examples:

$\vdash T \Rightarrow Eval\ env\ [1]\ (int\ 1)$

$\vdash Eval\ env\ [x]\ (int\ x) \Rightarrow Eval\ env\ [x]\ (int\ x)$

Automation

Examples:

$$\vdash \top \Rightarrow \text{Eval env } [1] \text{ (int 1)}$$

$$\vdash \text{Eval env } [x] \text{ (int } x) \Rightarrow \text{Eval env } [x] \text{ (int } x)$$

Lemmas ...

$$\begin{aligned} &\vdash \text{Eval env } x_1 \text{ (int } n_1) \Rightarrow \\ &\quad \text{Eval env } x_2 \text{ (int } n_2) \Rightarrow \\ &\quad \text{Eval env } [x_1 + x_2] \text{ (int } (n_1 + n_2)) \end{aligned}$$

... are used to prove compound terms:

$$\vdash \text{Eval env } [x] \text{ (int } x) \Rightarrow \text{Eval env } [x + 1] \text{ (int } (x + 1))$$

Closures are similar

Lemma for lambda:

$$\vdash (\forall v x. a x v \Rightarrow \text{Eval } (env [n \mapsto v]) \text{ body } (b (f x))) \Rightarrow \\ \text{Eval } env [\text{fn } n \Rightarrow \text{body}] ((a \longrightarrow b) f)$$

relation based on relations for input (a) and output (b)

Allows us to prove:

$$\vdash T \Rightarrow \text{Eval } env [\text{fn } x \Rightarrow x + 1] ((\text{int} \longrightarrow \text{int}) (\lambda x. x + 1))$$

... from result from previous slide, i.e.

$$\vdash \text{Eval } env [x] (\text{int } x) \Rightarrow \text{Eval } env [x + 1] (\text{int } (x + 1))$$

Mini demo

Implemented in HOL4 as a proof-producing tool

Easy to implement

Handles pure ML-like subset of higher-order logic

Scales to full CakeML compiler

Recent addition:

Support for monadic functions for I/O and stateful code

see IJCAR'18 preprint at cakeml.org

(Significantly improved speed of register allocator.)

Synthesis from monadic funs

For pure we use Eval:

$$\text{Eval } env \ exp \ post = \exists v. \langle env, exp \rangle \Downarrow v \wedge post \ v$$

For monadic functions we use EvalM:

$$\text{EvalM } ro \ env \ st \ exp \ P \ H \iff$$

$\forall s.$

$$\text{state_rel } H \ st \ s \Rightarrow$$

$$\exists s_2 \ res \ st_2 \ ck.$$

$$(\text{evaluate } (s \text{ with clock } := ck) \ env \ [exp] = (s_2, res)) \wedge$$

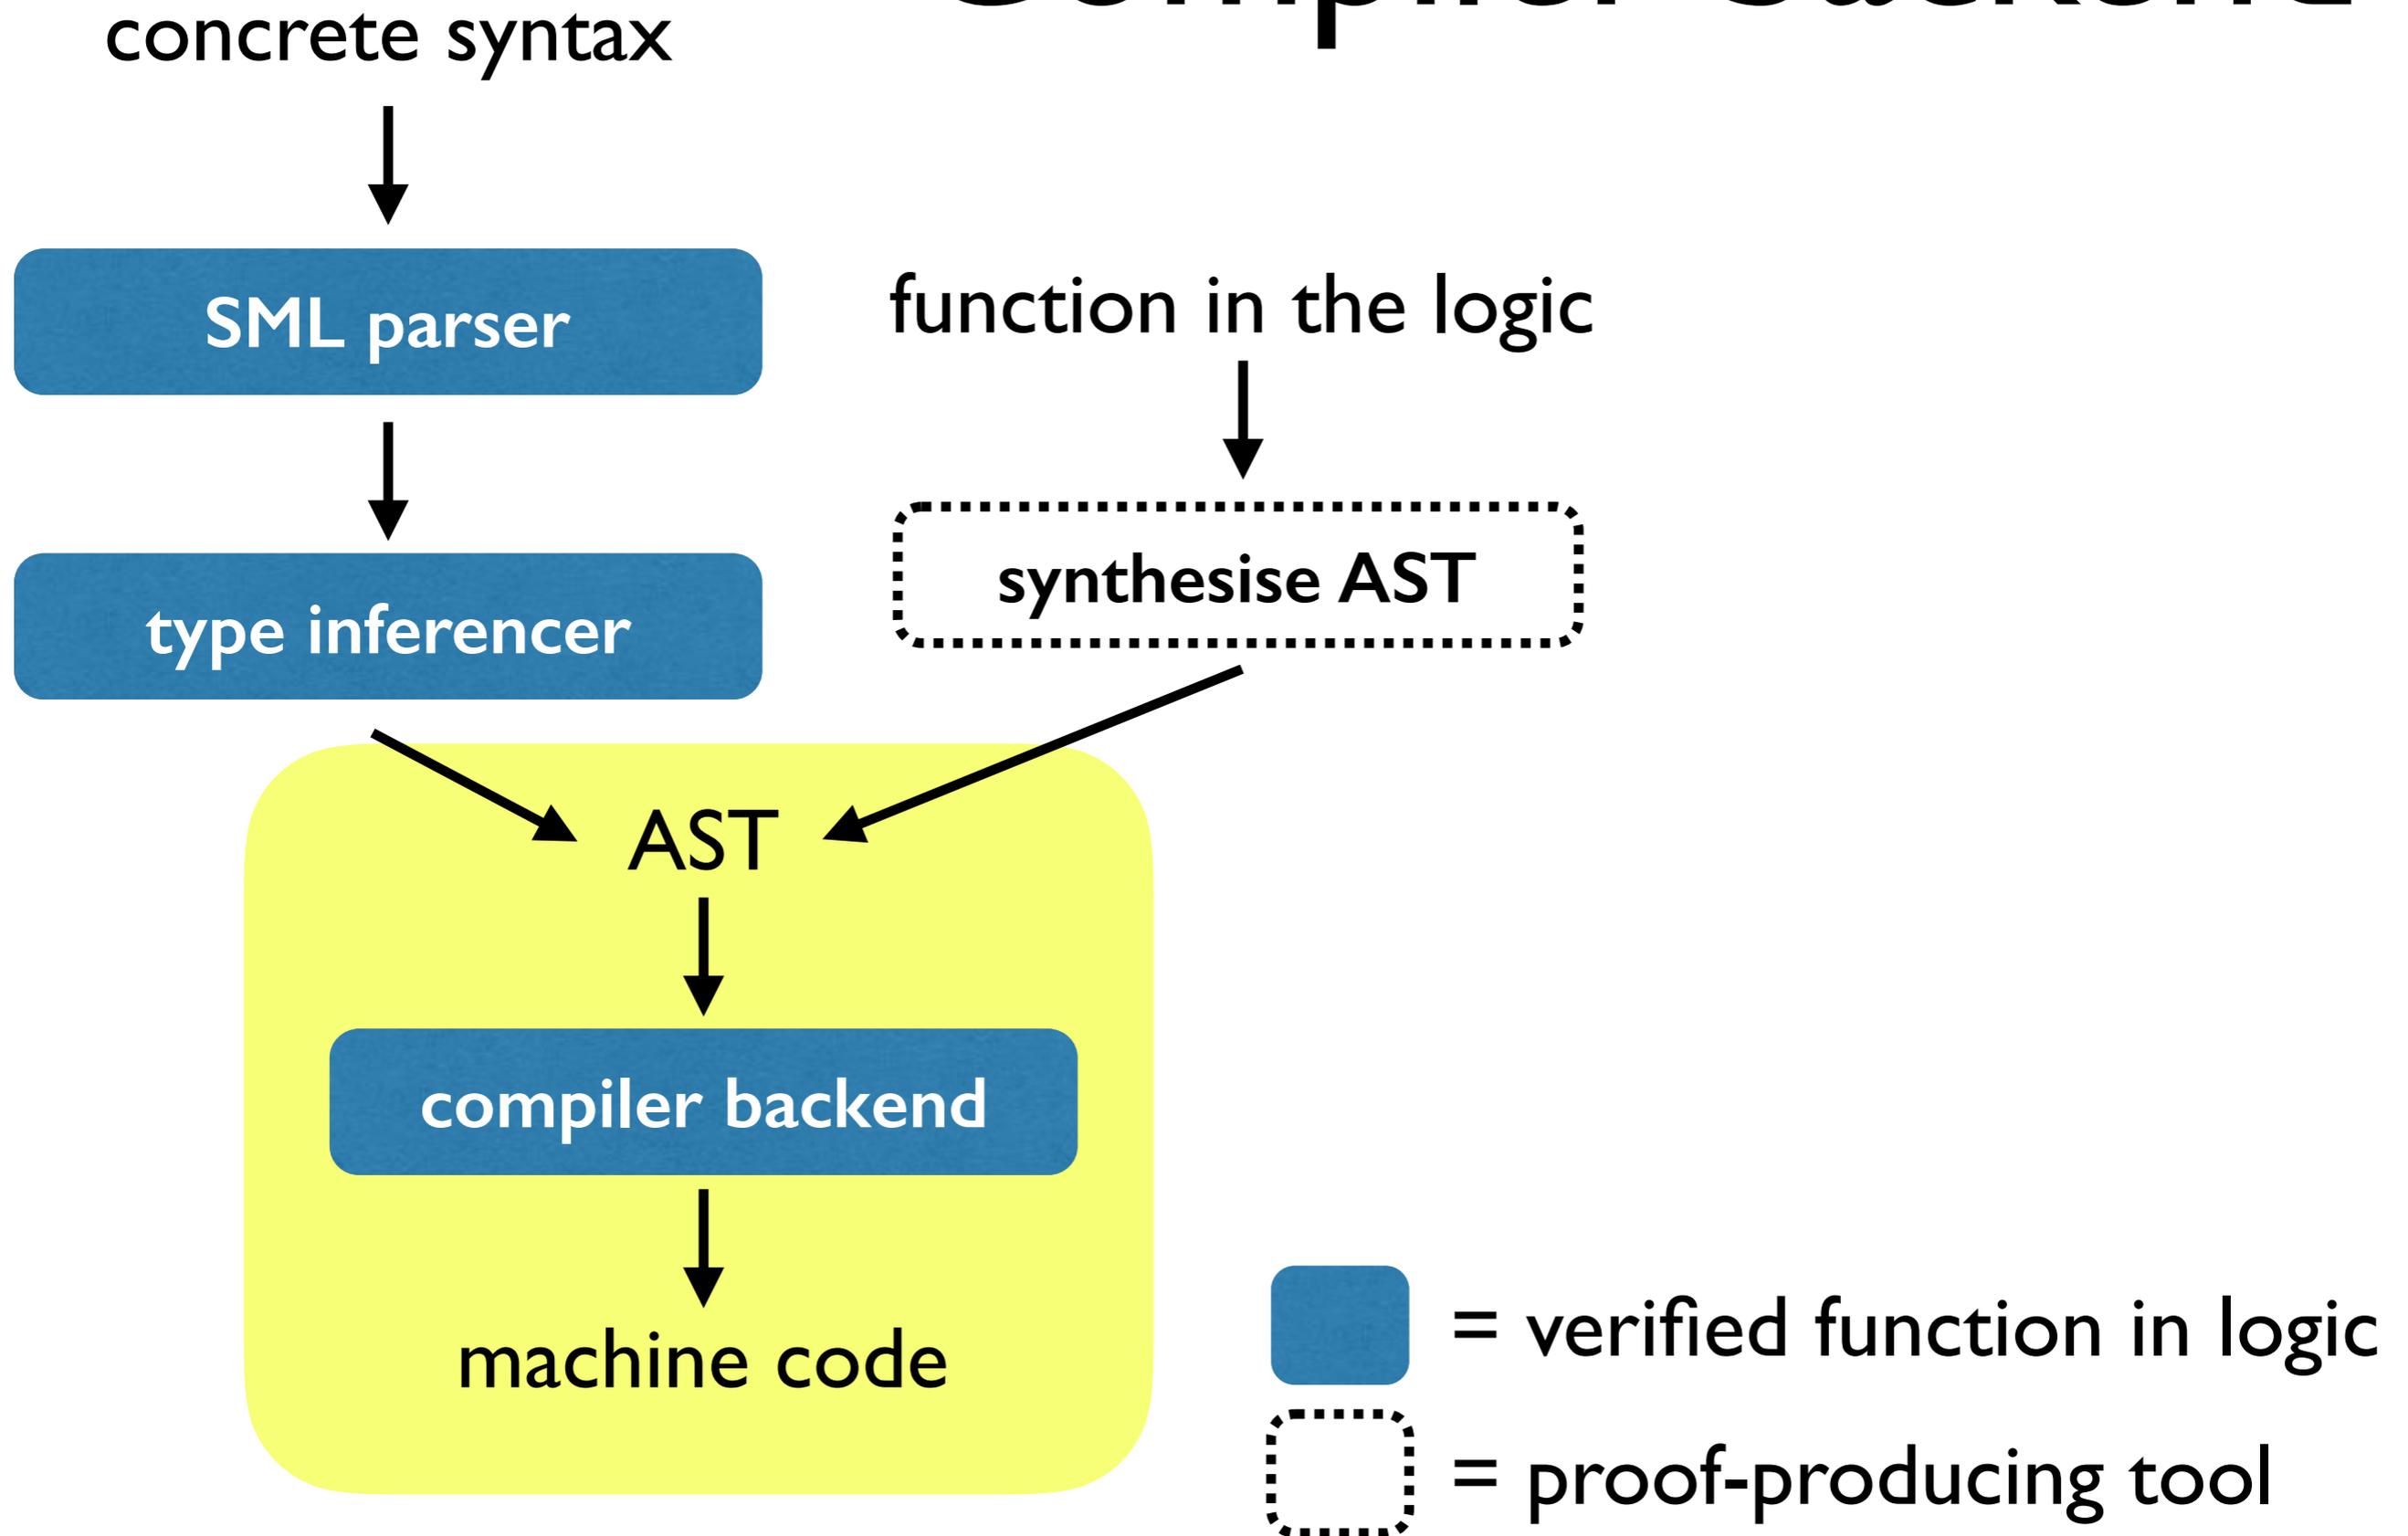
$$P \ st \ (st_2, res) \wedge \text{state_rel_frame } ro \ H \ (st, s) \ (st_2, s_2)$$

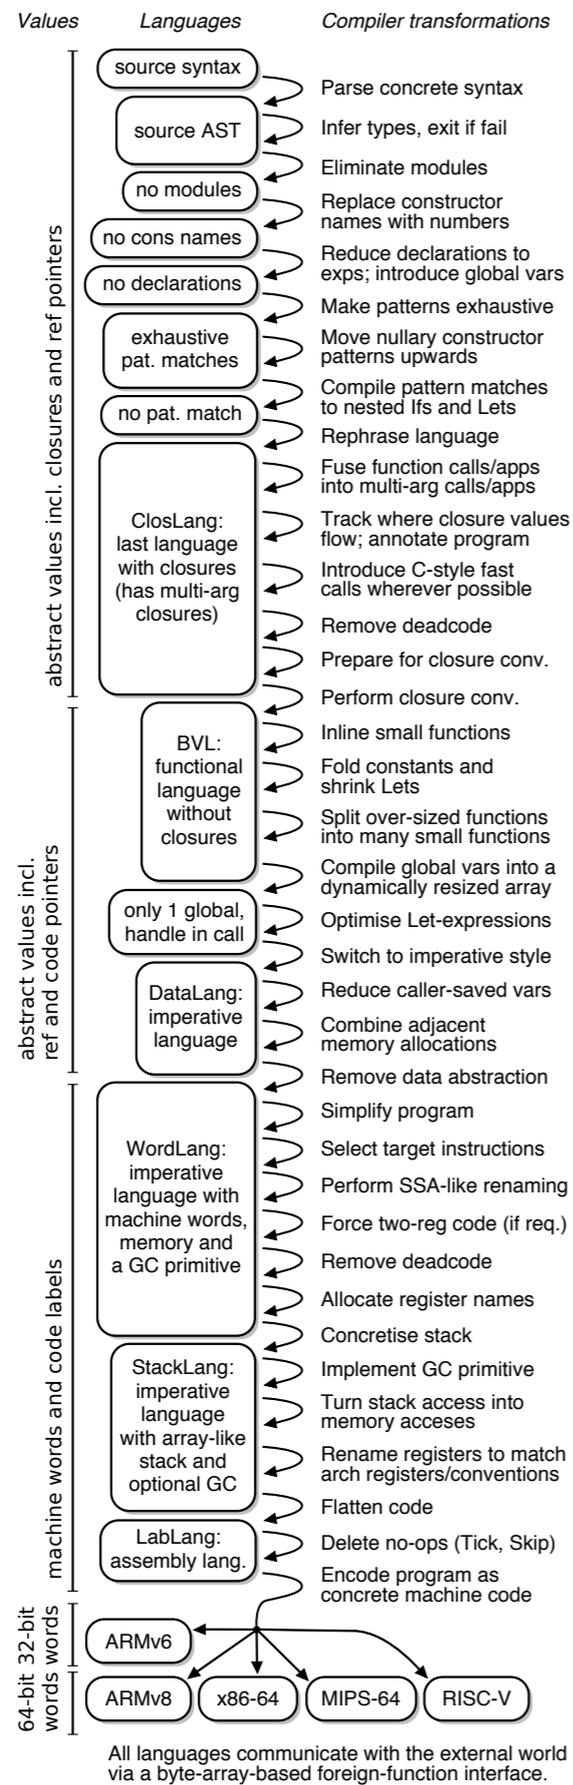
H is heap assertion relating monad state to CakeML state

where: $\text{state_rel } (h, p) \ st \ s \iff (h \ st \ * \ T) \ (\text{st2heap } p \ s)$

separating conjunction

Compiler backend





Compiler backend:

12 intermediate languages (ILs)
and many within-IL optimisations

each IL at the right level of abstraction

for the benefit of
 proofs and compiler
 implementation

Next slide zooms in

Values used by the semantics

Both proved sound and complete.

Values

Languages

Compiler transformations

source syntax

Parse concrete syntax

source AST

Infer types, exit if fail

no modules

Eliminate modules

no cons names

Replace constructor names with numbers

no declarations

Reduce declarations to exps; introduce global vars

exhaustive pat. matches

Make patterns exhaustive

Move nullary constructor patterns upwards

no pat. match

Compile pattern matches to nested ifs and Lets

Rephrase language

ClosLang:
last language with closures
(has multi-arg closures)

Fuse function calls/apps into multi-arg calls/apps

Track where closure values flow; annotate program

Introduce C-style fast calls wherever possible

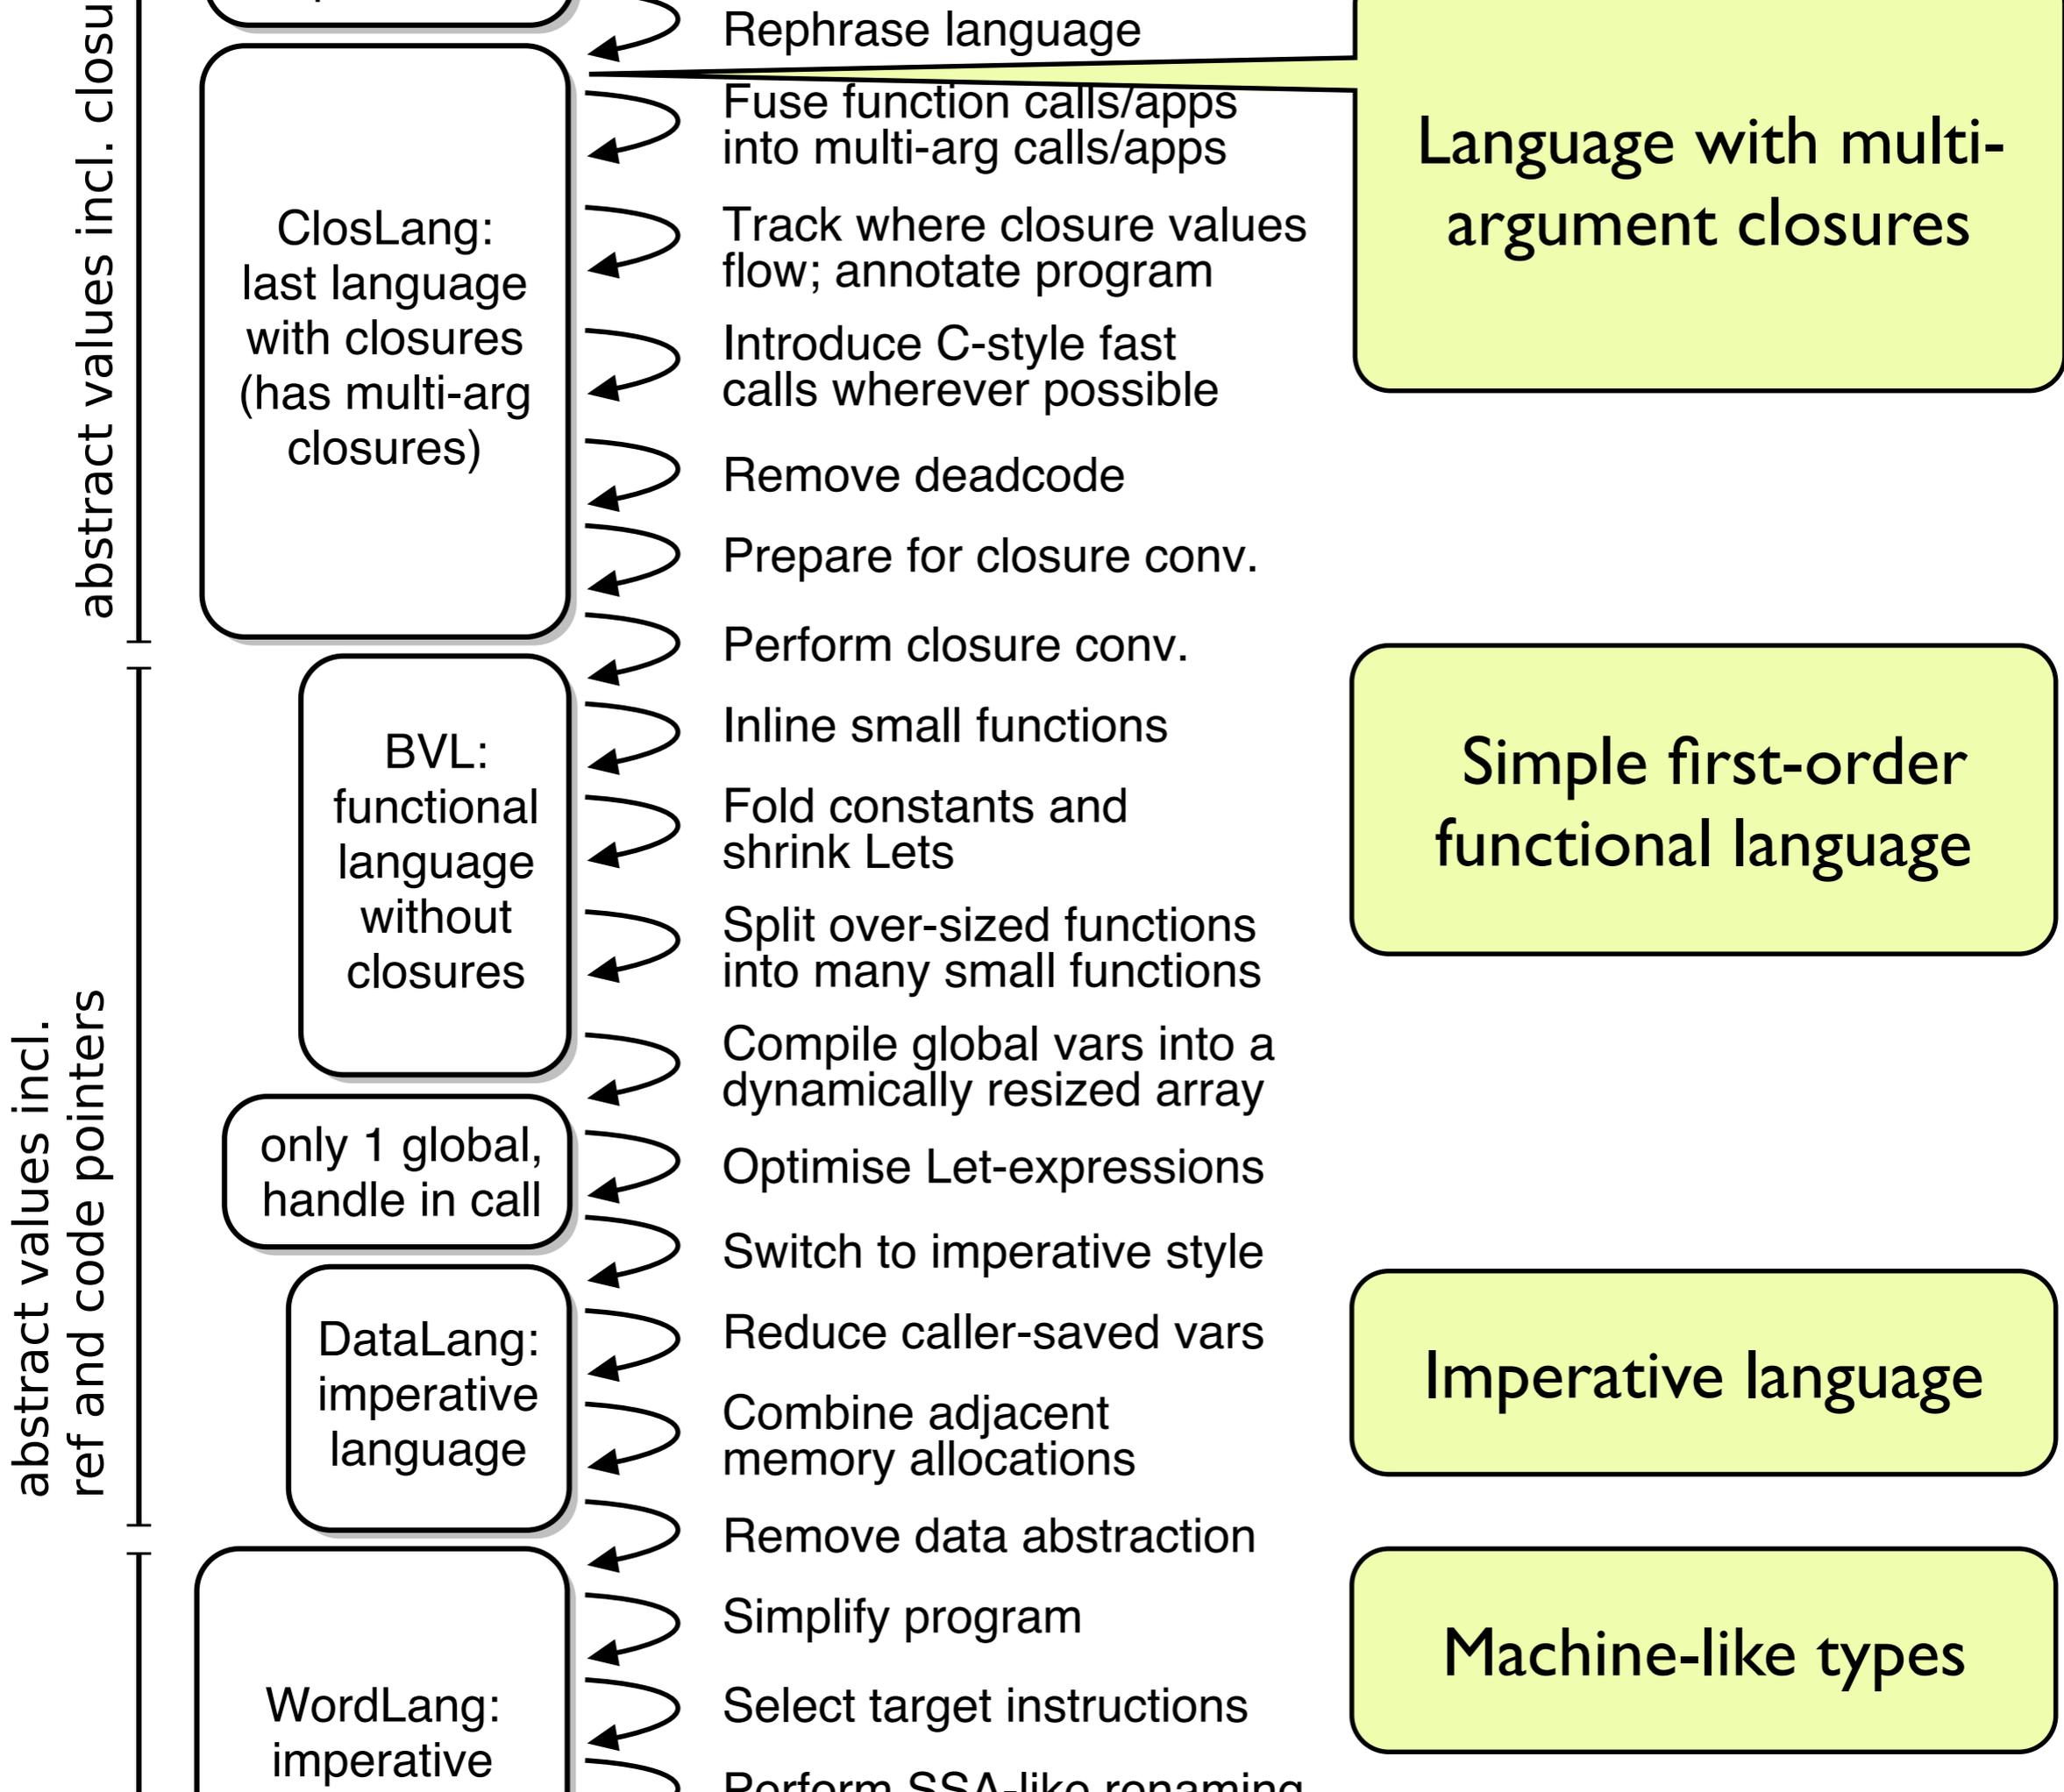
Remove deadcode

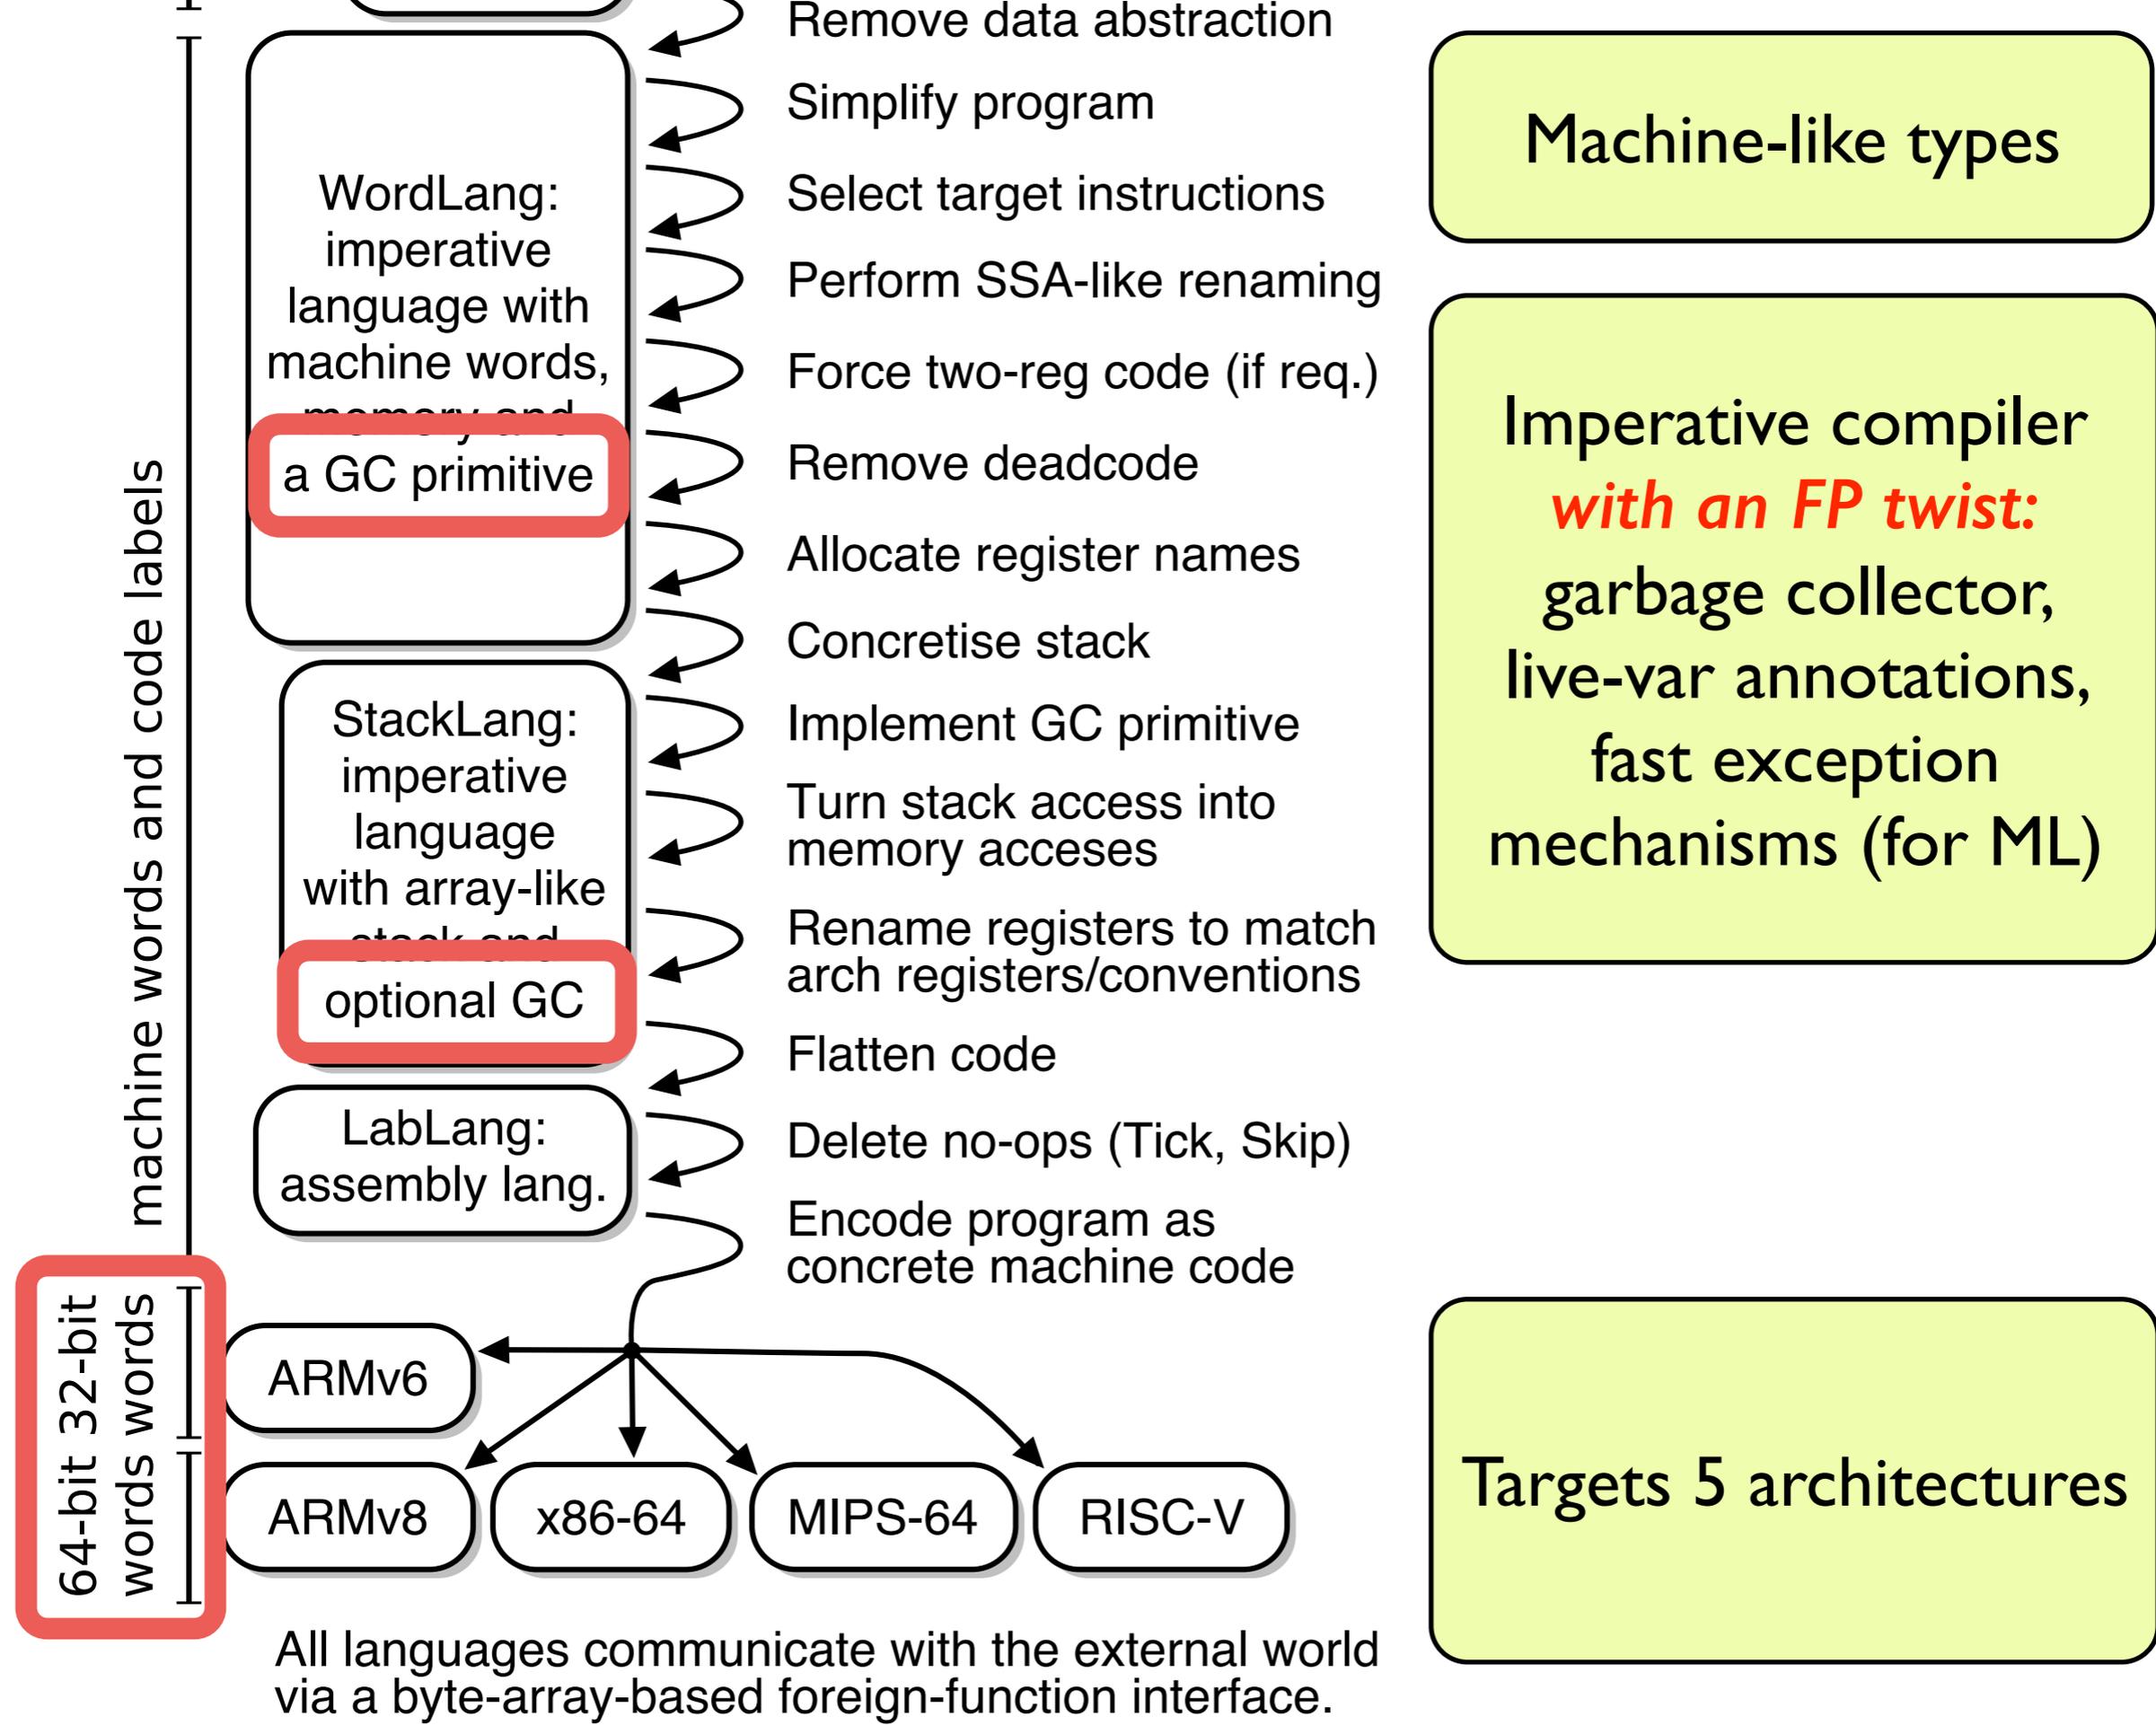
Parser and type inferencer as before

Early phases reduce the number of language features

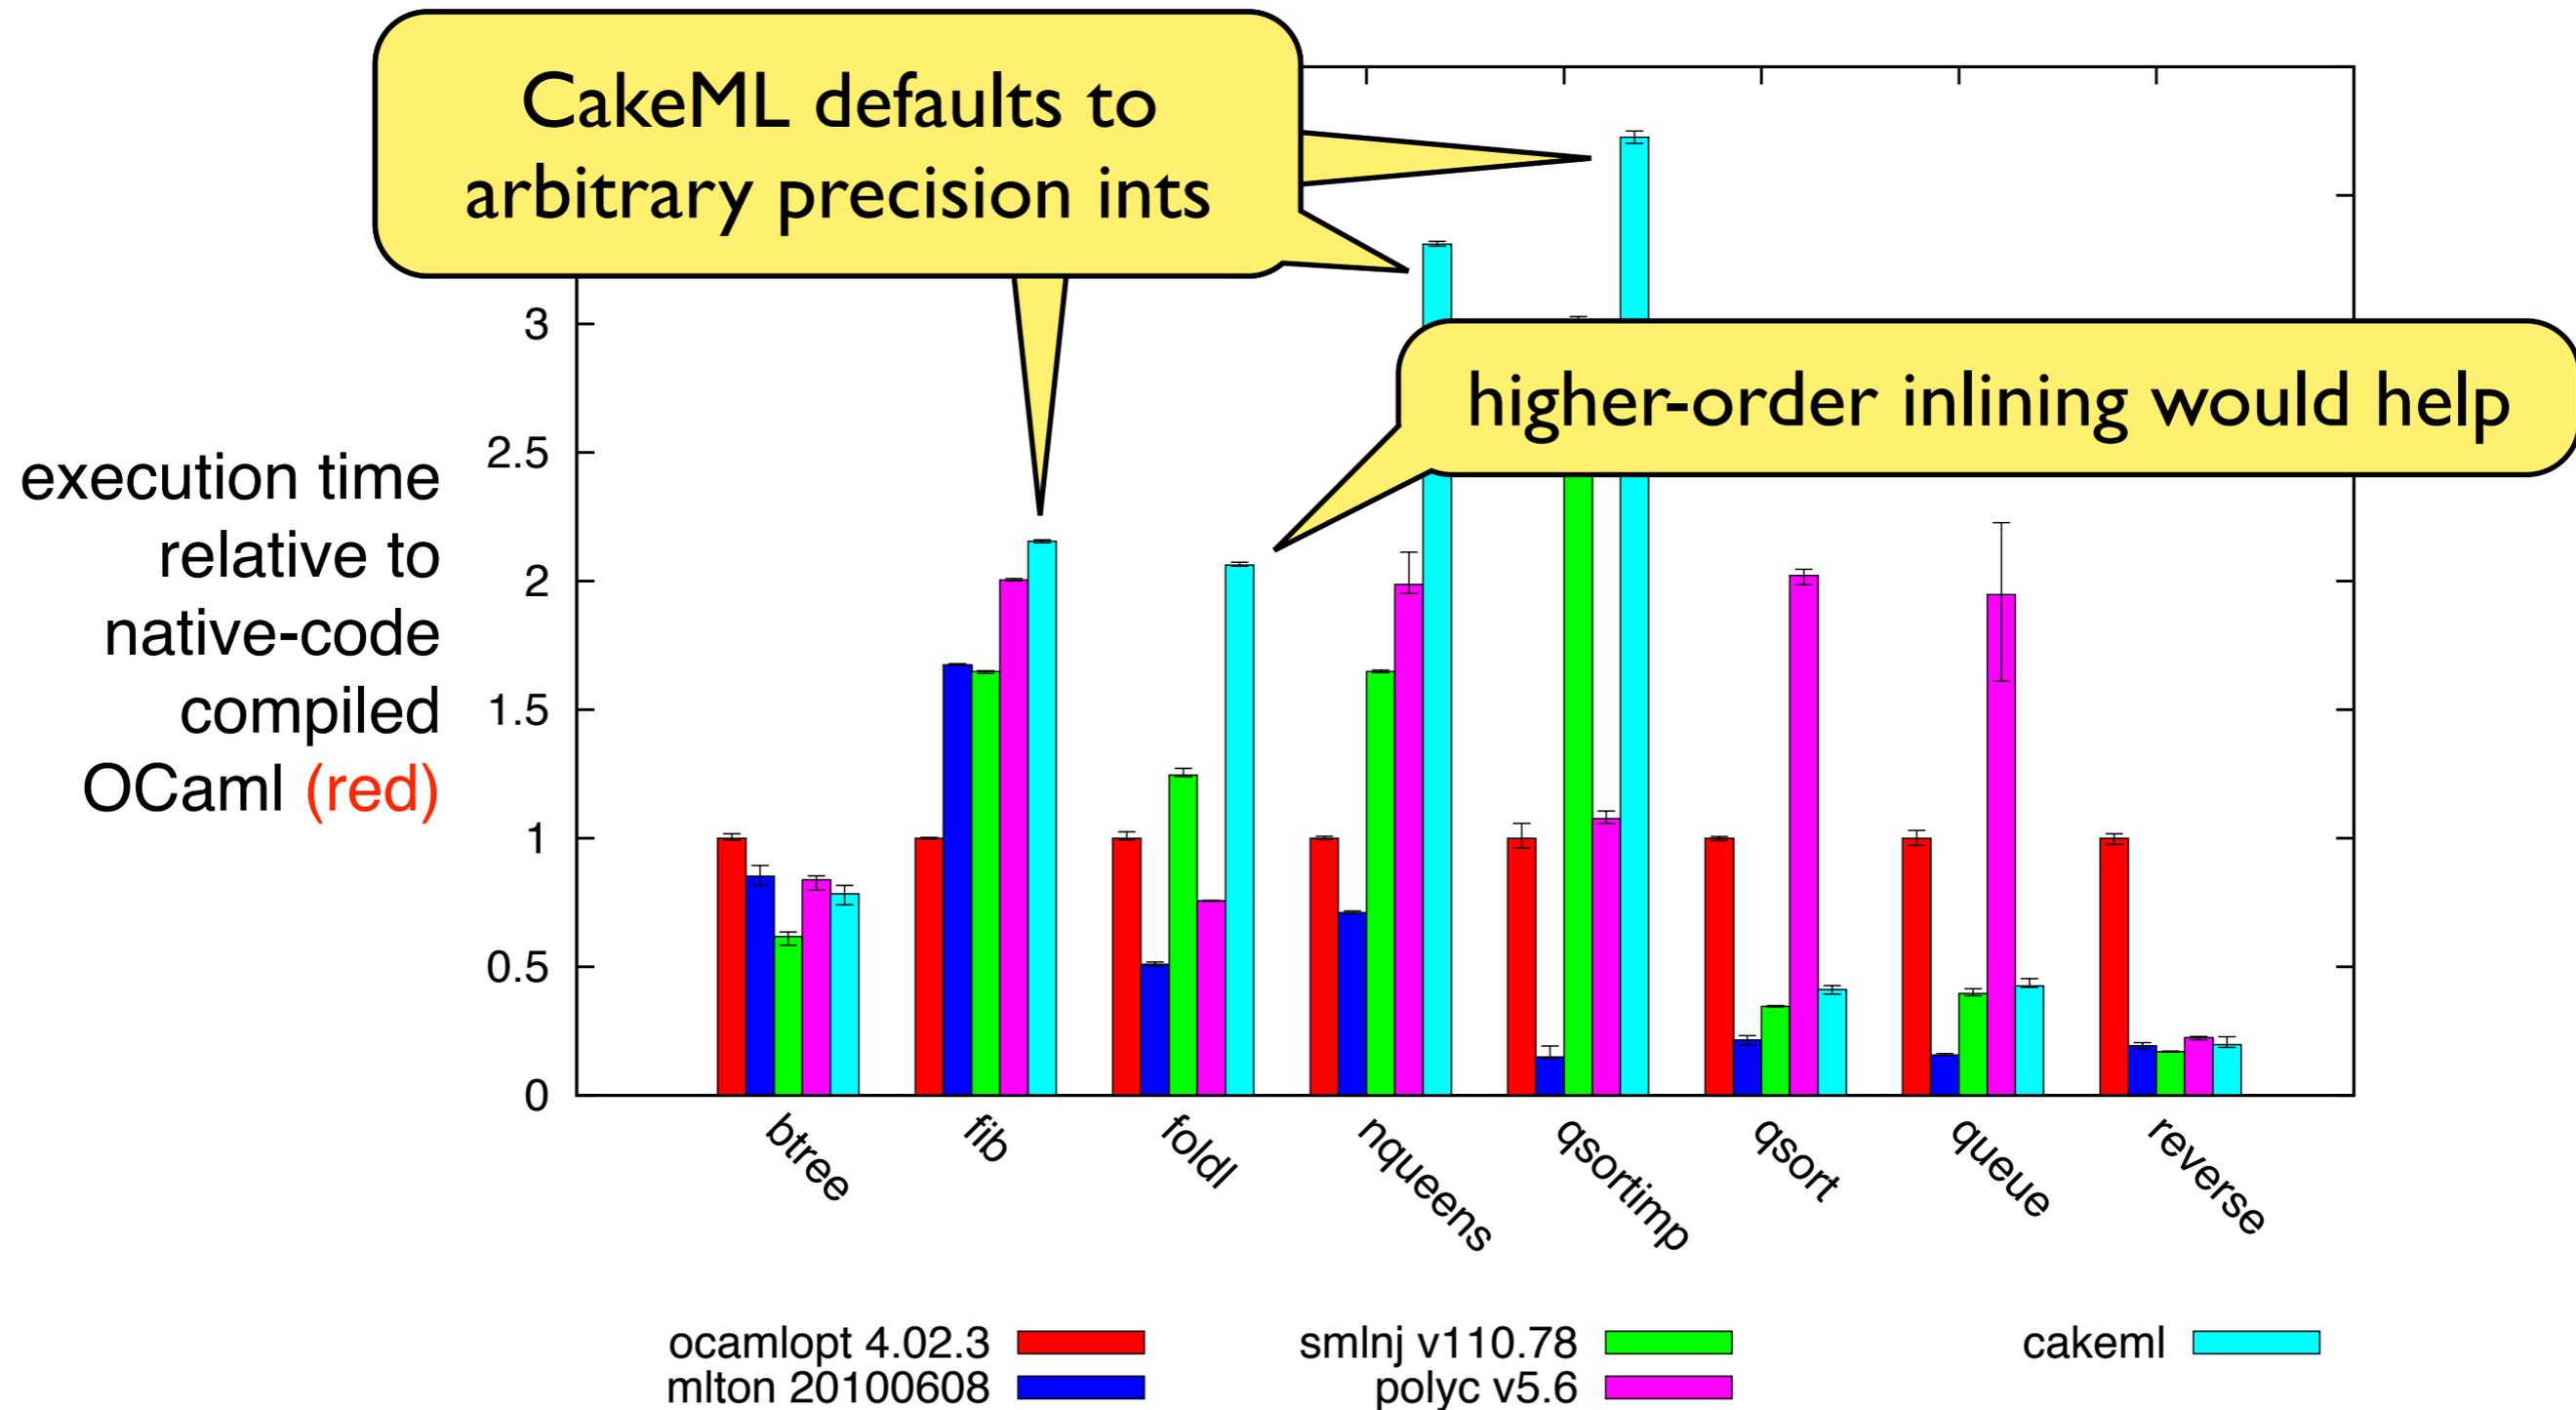
Language with multi-argument closures

Track values incl. closures and ref pointers





(Old) Performance numbers



Last slide

Questions?

Summary:

CakeML provides a *code extraction mechanism* that relates *via proof* the extracted (machine) code with the original functions from the logic of an ITP.

includes support for I/O and stateful code (monads)

In the pipeline:

- A major new language feature (and applications)
- Support for proving correctness of infinite executions

Get involved! → CakeML is an open project: cakeml.org

Extra: Compiler correctness

$\vdash \text{config_ok } cc \ mc \Rightarrow$
 case compile cc prelude input of
 Success ($bytes, ffi_limit$) \Rightarrow
 \exists behaviours.
 cakeml_semantics ffi prelude input =
 Execute behaviours \wedge
 $\forall ms.$
 code_installed ($bytes, cc, ffi, ffi_limit, mc, ms$) \Rightarrow
 machine_sem mc ffi $ms \subseteq$
 extend_with_resource_limit behaviours
 | Failure ParseError \Rightarrow
 cakeml_semantics ffi prelude input = CannotParse
 | Failure TypeError \Rightarrow
 cakeml_semantics ffi prelude input = IllTyped
 | Failure CompileError \Rightarrow true