

# PureCake

## A Verified Compiler for a Lazy Functional Language

Monday 19th June — PLDI 2023

---

**Hrutvik Kanabar** *University of Kent*

**Samuel Vivien** *Chalmers University of Technology, École Normale Supérieure PSL*

**Oskar Abrahamsson** *Chalmers University of Technology*

**Magnus O. Myreen** *Chalmers University of Technology*

**Michael Norrish** *Australian National University*

**Johannes Åman Pohjola** *University of New South Wales*

**Riccardo Zanetti** *Chalmers University of Technology*

## Implementing MyCriticalSoftware



?



type safety  
memory safety



type safety  
memory safety  
purity vs. I/O  
ref. transparency  
laziness  
free theorems

## Compiling MyCriticalSoftware



# Why verified compilation?

How do we know what our programs will do?

understand  
source semantics + trust compiler  
to preserve it

We may trust **GHC** — but does its source semantics match our understanding?

Verified compilation gives us a **verified link** between:

- **formally-specified source semantics**
- **semantics-preserving compilation**

The **PureCake** project:

a HOL4-**verified compiler** for  
a **lazy, purely functional** language which  
is inspired by **Haskell** and  
targets  **CakeML**

*CakeML = a verified implementation of a subset of ML [POPL14]*

## Highlights

- sound equational reasoning
- parsing expression grammar (PEG) for Haskell-like syntax
- two-phase constraint-based type inference\*
- demand analysis\*
- optimisations for non-strict idioms
- monadic reflection\* (monadic  $\rightarrow$  imperative)
- CakeML as a back end for end-to-end verified compilation

**\*Not mechanically verified before**

## High level, whistle-stop tour!

For more details:

- Read our paper: 📄 [cakeml.org/pldi23-purecake.pdf](https://cakeml.org/pldi23-purecake.pdf)
- Visit our GitHub: 🐙 [github.com/cakeml/pure](https://github.com/cakeml/pure)
- Talk to us!

Introduction

**Source language**

Compiler front end

Compiler back end

Connection with CakeML



# A realistic functional language

## PureLang has standard functional idioms ...

```
fact :: Integer -> Integer -> Integer
fact a n =
  if n < 2 then a
  else fact (a * n) (n - 1)
```

general recursion

```
map :: (a -> b) -> [a] -> [b]
map f l = case l of
  [] -> []
  h:t -> f h : map f t
```

algebraic data types +  
pattern-matching

```
factorials :: [Integer]
factorials = map (fact 1) (numbers 0)
```

higher-order functions

# A realistic subset of Haskell

## ... and Haskell extras

```
numbers :: Integer -> [Integer]
numbers n = n : numbers (n + 1)
```

```
main :: IO ()
main = do
  n <- readInt -- read from stdin
  let facts = take n factorials
  app (\i -> print $ toString i) facts
```

laziness  $\rightarrow$  infinite data


pure by default, monads for:

- sequencing
- stateful computations
- I/O

Single **IO** monad for arrays, exceptions, and I/O (via FFI calls)

**Also:** indentation-sensitivity, **do** notation, mutual recursion, ...

A tale of two ASTs... separate implementation and verification



- higher-level
  - used in implementation
  - includes **case**
- ground truth for semantics
  - constructor operations:  
test name/arity equality &  
argument projection

## Need to model non-termination and I/O

HOL4-expressible coinductive *interaction trees* [Xia et. al., POPL20]:

itree $E A R ::= \text{Ret } R$	<i>termination</i>
Vis $E (A \rightarrow \text{itree } E A R)$	<i>I/O via FFI channel</i>
Div	<i>silent divergence</i>



vs.



### No Tau nodes? Div!?

rely on non-constructivity of HOL4's logic  
*strong bisimulation coincides with equality*

## Operational semantics in layers:

1. Weak-head evaluation: call-by-name, functional big-step

$$\text{eval}_{\text{wh}}^n e = wh$$

2. Lift to unlocked evaluation

$$\text{eval}_{\text{wh}} e \stackrel{\text{def}}{=} \begin{cases} wh & \text{for **some** } n, \text{ eval}_{\text{wh}}^n e = wh \neq \text{Timeout} \\ \text{Timeout} & \text{for **all** } n, \text{ eval}_{\text{wh}}^n e = \text{Timeout} \end{cases}$$

3. Stateful interpretation of **IO** operations

$$\langle \! \langle - , - , - \rangle \! \rangle : wh \rightarrow \kappa \rightarrow \sigma \rightarrow \text{itree } E A R$$

**Finally,**  $\llbracket e \rrbracket \stackrel{\text{def}}{=} \langle \! \langle \text{eval}_{\text{wh}} e, \varepsilon, \emptyset \rangle \! \rangle$

# Mechanised equational reasoning

**Mechanise untyped applicative bisimulation**,  $\cong \cong$  [Abramsky, 1990]

Proved *congruent* via Howe's method [Howe, 1996]

*i.e. bisimilar sub-expressions  $\implies$  bisimilarity*

## Definitions:

- $\alpha$ -equivalence:  $e_1 =_\alpha e_2 \stackrel{\text{def}}{=} \text{perm\_vars } e_1 e_2$
- $\beta$ -equivalence:  $(\lambda x. e_1) \cdot e_2 =_\beta (\text{freshen}_{e_2} e_1)[e_2/x]$
- A standard contextual equivalence:  $e_1 \sim e_2$   
(*equality under all closing contexts*)

## Derived results:

$$e_1 \cong e_2 \iff e_1 \sim e_2 \qquad \frac{e_1 =_\alpha e_2}{e_1 \cong e_2} \qquad \frac{e_1 =_\beta e_2}{e_1 \cong e_2}$$

Standard Hindley-Milner rules... with an unusual soundness proof

## Problem:

non-strict semantics + exhaustive **case** splits mean that  
“preservation” (subject reduction) does not hold

## Solution:

- Define an alternative syntax for typing
- Prove subject reduction by construction
- Use equational theory to bridge the gap to original syntax

Introduction

Source language

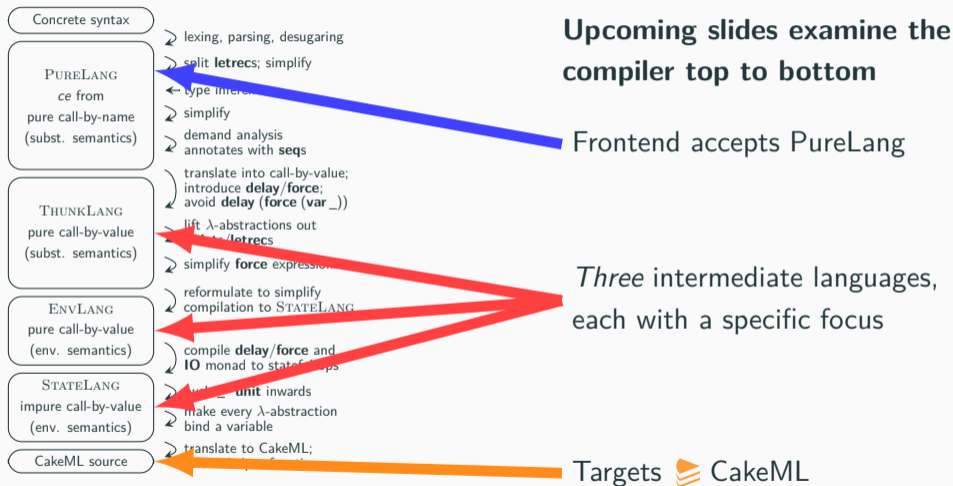
**Compiler front end**

Compiler back end

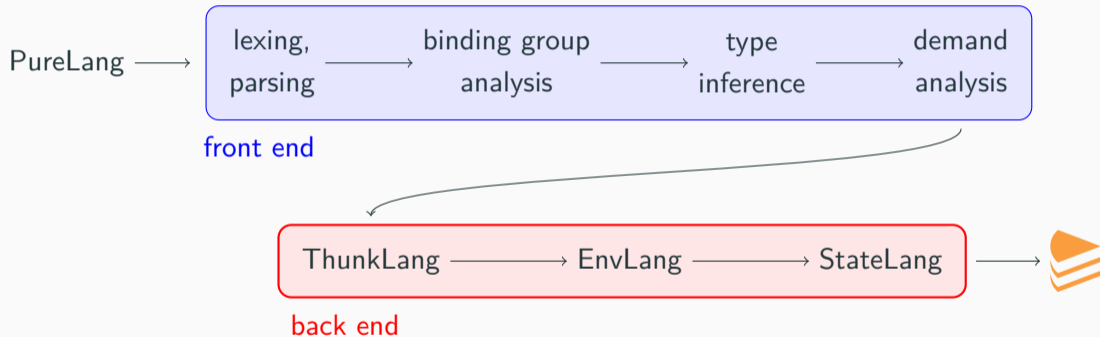
Connection with CakeML



# Compiler structure



# Compiler structure



## Indentation-sensitive parsing expression grammar (PEG):

 PEG + [Adams POPL13]

$$N \rightarrow X_1^{\mathcal{R}_1} X_2^{\mathcal{R}_2} \dots X_n^{\mathcal{R}_n}$$

where  $\mathcal{R} \in \{=, >, \geq, \mathcal{U}\}$

- Symbolic sets of possible indentations for each non-terminal
- Verified to terminate on all inputs

# Binding group analysis

Parsing

```
z = 42
y = x + 1
x = w + y
w = 0
main
```

→

```
let rec z = 42
    y = x + 1
    x = w + y
    w = 0
in main
```

Analyse dependencies



*Pseudo*-topological sort



Transform code + tidy

```
let w = 0 in
let rec x = w + y ; y = x + 1
in main
```

**Verified entirely within equational theory**

# Sound constraint-based type inference

**Two-phases:** generate *all* constraints  $\longrightarrow$  solve constraints

Subset of Helium's  $\text{TOP}$  framework [Heeren et al., Haskell 2003]

- Open to high-quality error messages
- Path to various Haskell 98 features

**Soundness theorem:**

infer  $ce$  **succeeds**

$\implies ce \vdash_{\text{TOP}} cs$  *and*  $cs$  **solveable**

$\implies \Gamma \vdash ce : \tau$

**Avoid excessive thunks** — acc heap-allocated each recursive call!

```
fact acc n =  
  if ... then acc else fact (acc * n) (n - 1)
```

- $e$  **demands**  $\overline{x_n}$   $\stackrel{\text{def}}{=} e \cong (x_1 \text{ `seq` } \dots x_n \text{ `seq` } e)$
- Implement/verify\* analysis:  $e$  **demands** (analyse  $e$ )
- Prefix code with `seq`, including in recursive functions

\*Verify with an alternative equational theory,  $\approx$  [Sergey et al., 2014]:

stuck  $\approx$  diverged  $\approx \perp$     **but**    well-typed  $\implies \approx, \cong$  coincide  
seq-prefixing preserves typing

Introduction

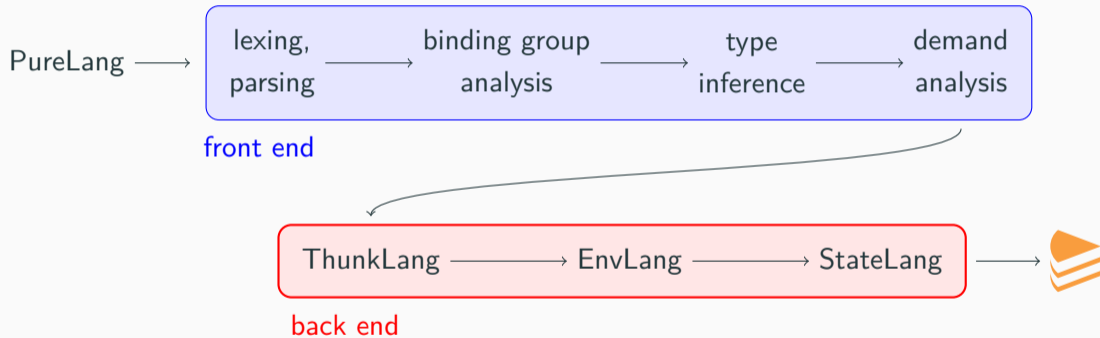
Source language

Compiler front end

**Compiler back end**

Connection with CakeML

# Compiler structure





# Methodology — implementation vs. verification

**Prior work:** (such as CakeML)

- Define implementation function:  $\text{transform} : e \rightarrow e$
- Verify:  $\text{wf } e \implies \llbracket \text{transform } e \rrbracket = \llbracket e \rrbracket$

**This work:**

$$e \mathcal{R} e'$$

*syntactic relations*

- for **verification**
- an implementation *envelope*

$$\text{compile } ce = ce'$$

*code transformation*

- for **implementation**
- must fit in relation envelope

1. **Define** and **verify**  $\mathcal{R}$ :  $e \mathcal{R} e' \implies \llbracket e \rrbracket = \llbracket e' \rrbracket$   
*Three simulation proofs: one per layer of the semantics*
2. **Define** `compile` :  $ce \rightarrow ce$
3. **Verify**  $wf\ ce \implies (\text{desugar } ce) \mathcal{R} (\text{desugar } (\text{compile } ce))$
4. **Compose** theorems:  
 $wf\ ce \implies \llbracket \text{desugar } ce \rrbracket = \llbracket \text{desugar } (\text{compile } ce) \rrbracket$
5. **Integrate** into compiler, discharge  $wf\ ce$

**Separation of concerns for modularity and ease-of-verification**

## Call-by-value semantics

*Syntax:*  $e ::= \dots \mid \mathbf{delay} e \mid \mathbf{force} e$  introduce *thunks*

*Semantics:*

$$\text{eval}(\mathbf{delay} e) = \mathbf{thunk} e$$
$$\frac{\text{eval } e' = v}{\text{eval}(\mathbf{force} e) = v}$$

eval  $e = \mathbf{thunk} e'$

*NB* thunks are pure, value-sharing comes later

*Optimisation:* reduce  $\mathbf{delay}(\mathbf{force} x)$ ; two forms of restricted CSE

*Verification:* seven syntactic relations in total

## Environment-based semantics + minor reformulations

*Syntax:* essentially unchanged

*Semantics:* ~~substitutions~~ **environments**

*Verification:* focuses on the change in semantic style


IO monad compiled to **effectful primitives, thunks shared statefully**

*Syntax:*  $e ::= \dots \mid \text{malloc } n \ e \mid \dots$  remove **delay/force**,  
**return/bind/...**

*Semantics:* *stateful* CESK machine

*Compilation:*  $\text{alloc } n \ e \mapsto \lambda_. \text{malloc } n \ e$   
 $\text{force } e \mapsto \text{let } x = e' \text{ in}$   
 $\text{if } x[0] \text{ then } x[1]$   
 $\text{else } \dots x[0] := \text{true}; x[1] := v \dots$

[true, v] or  
[false, λ\_. ...]



*Optimisation:* simplify  $\lambda_. e$  and **unit**

# Roadmap

Introduction

Source language

Compiler front end


Compiler back end

**Connection with CakeML**

## Reconciling differing semantic styles

$$o_1 \xrightarrow{\Delta(o_1)} o_2 \xrightarrow{\Delta(o_2)} \dots$$

**linear** oracles:  $\text{semantics}_{\Delta} e = tr$



**branching** ITrees:  $\llbracket e \rrbracket = \text{Vis } \dots$

- Verified ITree semantics:  $\llbracket e \rrbracket \Rightarrow \mathcal{A} tr \Leftrightarrow \text{semantics}_{\Delta} e = tr$
- New compiler correctness:

$$\frac{\text{cakeml } e = \text{Some } \textit{code}}{\llbracket \textit{machine} \rrbracket_M \text{ prunes } \llbracket e \rrbracket \Rightarrow}$$

$\text{purecake } str = \text{Some } ast \Rightarrow$

---

**exists**  $ce$  **such that**

$\text{frontend } str = \text{Some } (ce, \_)$

**$ce$  is type safe**

$\llbracket \text{desugar } ce \rrbracket_{\text{pure}} \simeq \llbracket ast \rrbracket \Rightarrow \Rightarrow$



$\text{purecake } str = \text{Some } ast$

$\text{cakeml } ast = \text{Some } code$

*code in memory of machine*

---

**exists**  $ce$  **such that**

$\text{frontend } str = \text{Some } (ce, -)$

$\llbracket machine \rrbracket_M \text{ prunes } \llbracket \text{desugar } ce \rrbracket_{\text{pure}}$

**Verified bootstrapping** using *proof-producing synthesis* [ICFP12]

HOL4 functions  $\xrightarrow{\text{synthesise}}$  *AST* + *equivalence proof*

The PureCake compiler is a HOL4 function...

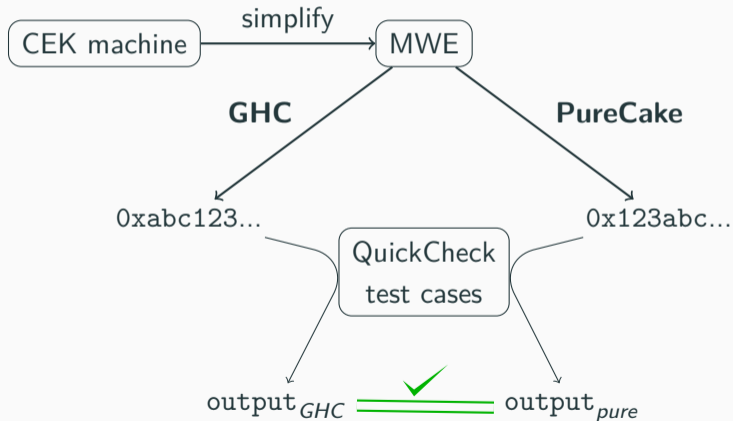
*purecake*  $\xrightarrow{\text{synthesise}}$  *purecake\_AST* + *equivalence proof*

... and the CakeML compiler can be evaluated in-logic:

$\vdash \text{cakeml} (\text{purecake\_AST}) = \text{0xabc123...}$

**Equivalence proofs transport verification to the binary**

## Testing on the Cardano block chain platform by QuviQ



## PureCake

 [cakeml.org/purecake](https://cakeml.org/purecake)

a **verified compiler** for a Haskell-like language

- sound equational reasoning
- Haskell-like syntax
- two-phase constraint-based sound type inference
- verified demand analysis
- optimisations to handle non-strict code realistically
- end-to-end guarantees by targeting CakeML
- feasible to use on real code

## Questions?

**Backup slides**

**Only a first version!** Many possible extensions, for example:


- Increasing source **expressivity** (e.g. for **case**)
- More Haskell 98 types, e.g. **typeclasses**
- More effective **demand analysis**
- Back end **optimisations**

A **verified REPL** for PureCake [*Sewell et. al., PLDI23*]

### Measure **execution time** and **memory allocations**

- Turn off individual optimisations to highlight their effect
  - pure: binding group analysis
  - demands: demand analysis
  - thunk: some **force** (**delay**  $e$ ) reduction and CSE in ThunkLang
  - state:  $\lambda$ -.  $e$ /**unit** optimisations in StateLang
- Five benchmarks, each accepting integer  $n$  input
  - primes:  $n$ th prime calculation
  - collatz: longest Collatz sequence for a number less than  $n$
  - life: Conway's Game of Life for  $n$  generations
  - queens: solutions for the  $n$ -queens problem
  - qsort: imperative quicksort for an array of length  $n$

# Evaluation — results



## Results

- ThunkLang optimisations provide significant benefit
- StateLang optimisations improve monadic code particularly
- Binding group analysis negligible
- Regressions for demand analysis: `seq`-insertion is too eager!



## Alternative compiler correctness

frontend  $str = \text{Some } (ce, \_)$

---

**$ce$  is type safe**

**exists  $ast$  such that**

purecake  $str = \text{Some } ast$

$\llbracket \text{desugar } ce \rrbracket_{\text{pure}} \approx \llbracket ast \rrbracket$