

# Verified compilers

Magnus Myreen

Chalmers University of Technology

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

# Verified compilers

*What?*

- Comes with a machine-checked proof that for any program, which does not generate a compilation error, the source and target programs behave identically

(Sometimes called *certified* compilers, but that's misleading...)

## Your program crashes.

Where do you look for the fault?

- Do you look at *your source code*?
- Do *look at the code for the compiler that you used*?

users want to rely on compilers

### Trusting the compiler

#### Bugs

When finding a bug, we go to great lengths to find it in our own code.

- Most programmers trust the compiler to generate correct code
- The most important task of the compiler is to generate correct code

Maybe it is worth the cost?

### Establishing compiler correctness

Cost reduction?

#### Alternatives

- Proving the correctness of a compiler is prohibitively expensive
- Testing is the only viable option

... but with testing you never know you caught all bugs!

## All (unverified) compilers have bugs

“Every compiler we tested was found to crash and also to silently generate wrong code when presented with valid input.”

PLDI'11

### Finding and Understanding Bugs in C Compilers

Xuejun Yang Yang Chen Eric Eide John Regehr

“[The verified part of] CompCert is the only compiler we have tested for which Csmith cannot find wrong-code errors. This is not for lack of trying: we have devoted about six CPU-years to the task.”

## Motivations

Bugs in compilers are not tolerated by users

Bugs can be hard to find by testing

Verified compilers must be used for verification of source-level programs to imply guarantees at the level of verified machine code

**Research question:** how easy (cheap) can we make compiler verification?

## This lecture: Verified compilers

**What?** Prove that compiler produces good code.

**Why?** To avoid bugs, to avoid testing.

**How?** By mathematical proof...

rest of  
this lecture

## Proving a compiler correct

### Ingredients:

- a **formal logic** for the proofs
- **accurate models** of
  - the **source** language
  - the **target** language
  - the **compiler** algorithm

like first-order logic, or higher-order logic

proofs are only about things that live within the logic, i.e. we need to represent the relevant artefacts in the logic

### Tools:

- a **proof assistant** (software)

a lot of details... (to get wrong)

... necessary to use mechanised proof assistant (think 'Eclipse for logic') to avoid accidentally skipping details

## Accurate model of prog. language

### Model of programs:

- syntax — what it looks like
- semantics — how it behaves

e.g. an **interpreter** for the syntax

### Major styles of (operational, relational) semantics:

- big-step — this style for structured source semantics
- small-step — this style for unstructured target semantics

... next slides provide examples.

## Syntax

### Source:

```
exp = Num num
      | Var name
      | Plus exp exp
```

### Target 'machine code':

```
inst = Const name num
       | Move name name
       | Add name name name
```

Target program consists of list of inst

## Source semantics (big-step)

Big-step semantics as **relation**  $\downarrow$  defined by **rules**, e.g.

$$\frac{}{(\text{Num } n, \text{ env}) \downarrow n} \quad \frac{\text{lookup } s \text{ in env finds } v}{(\text{Var } s, \text{ env}) \downarrow v}$$

$$\frac{(\text{x1}, \text{ env}) \downarrow v1 \quad (\text{x2}, \text{ env}) \downarrow v2}{(\text{Add x1 x2}, \text{ env}) \downarrow v1 + v2}$$

called "big-step": each step  $\downarrow$  describes complete evaluation

## Source semantics (...gone wrong)

Real-world semantics are not always clean:

<https://www.destroyallsoftware.com/talks/wat>

# Target semantics (small-step)

“small-step”: transitions describe parts of executions

We model the state as a **mapping from names to values** here.

```
step (Const s n) state = state[s ↦ n]
step (Move s1 s2) state = state[s1 ↦ state s2]
step (Add s1 s2 s3) state = state[s1 ↦ state s2 + state s3]

steps [] state = state
steps (x::xs) state = steps xs (step x state)
```

```
steps [] state = state
steps (x::xs) state = steps xs (step x state)
```

# Compiler function

The diagram illustrates the compiler's function through three examples of code generation, each with a callout explaining a specific feature:

- `compile (Num k) n = [Const n k]` → Callout: generated code stores result in register name (n) given to compiler
- `compile (Var v) n = [Move n v]` → Callout: Relies on variable names in source to match variables names in target.
- `compile (Plus x1 x2) n =  
 compile x1 n ++ compile x2 (n+1) ++ [Add n n (n+1)]` → Callout: Uses names above n as temporaries.

- Relies on variable names in source to match variables names in target.

# Correctness statement

*Proved using proof assistant — demo!*

For every evaluation in the source ...

for target state and k, such that ...

$\forall x \text{ env res.}$   
 $(x, \text{env}) \downarrow \text{res} \Rightarrow$   
 $\forall \text{state } k.$   
 $(\forall i \text{ v. } (\text{lookup env } i = \text{SOME } v) \Rightarrow (\text{state } i = v) \wedge i < k) \Rightarrow$   
 $(\text{let state}' = \text{steps } (\text{compile } x \text{ } k) \text{ state in}$   
 $(\text{state}' \text{ } k = \text{res}) \wedge$   
 $\forall i. i < k \Rightarrow (\text{state}' \text{ } i = \text{state } i))$

k greater than all var names and state in sync with source env ...

... in that case, the result res will be stored at location k in the target state after execution

... and lower part of state left untouched.

... and lower part of state left untouched.

[illegible]

*Well, that example was simple enough...*

But:

**Some people say:**  
A programming language isn't real until it has a self-hosting compiler

*Bootstrapping for verified compilers? **Yes!***

**POPL 2014**

# CakeML: A Verified Implementation of ML

Ramana Kumar<sup>1</sup> Magnus O. Myreen<sup>1</sup> Michael Norrish<sup>2</sup> Scott Owens<sup>3</sup>

<sup>1</sup> Computer Laboratory, University of Cambridge, UK  
<sup>2</sup> Canberra Research Lab, NICTA, Australia<sup>1</sup>  
<sup>3</sup> 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 (REPL) in x86-64 machine code. Our correctness theorem ensures that this REPL implementation prints only those results permitted by the semantics of CakeML. Our verification effort touches on a breadth of topics including lexing, parsing, type checking, incremental and dynamic compilation, garbage collection, arbitrary-precision arithmetic, and compiler bootstrapping.

Our contributions are twofold. The first is simply in presenting a system that is end-to-end verified, demonstrating that each piece of a verification effort can in practice be composed of a single proof. The second is in showing that the model ap-

## 1. Introduction

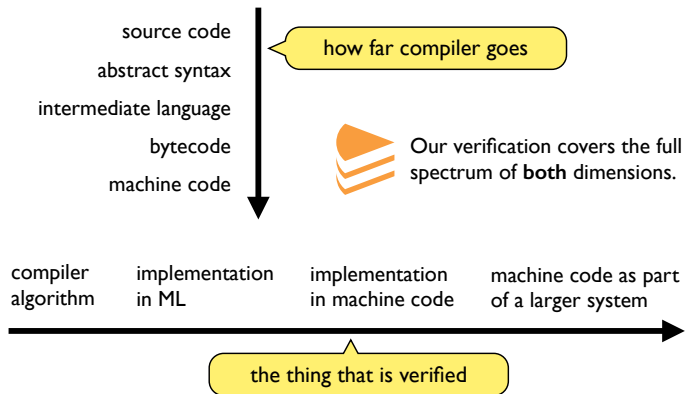
The last decade has seen a strong interest in verified compilation; and there have been significant, high-profile results, many based on the CompCert compiler for C [1, 14, 16, 29]. This interest is easy to justify: in the context of program verification, an unverified compiler forms a large and complex 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 of the following dimensions: one, the compilation of a compiler along two dimensions: one, the compilation of a compiler from a source language to a target language; and two, the compilation of a compiler from a source language to a target language.

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

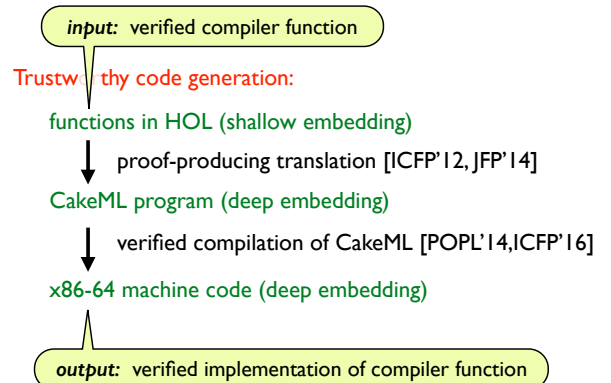
Practical CakeML and it is a source-to-source compiler that takes CakeML and OCaml code and produces OCaml code.

practicing CakeML, and it is a strong 2.0 ML and OCaml. By removing the code along-

# Dimensions of Compiler Verification



## Idea behind in-logic bootstrapping



## The CakeML at a glance

The CakeML language = Standard ML without I/O or functors

strict impure functional language

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

### The verified machine-code implementation:

parsing, type inference, compilation, garbage collection, bignums etc.


*implements a read-eval-print loop (see demo).*

## The CakeML *compiler verification*

## How?

Mostly **standard verification techniques** as presented in this lecture, but **scaled up** to large examples. (Four people, two years.)

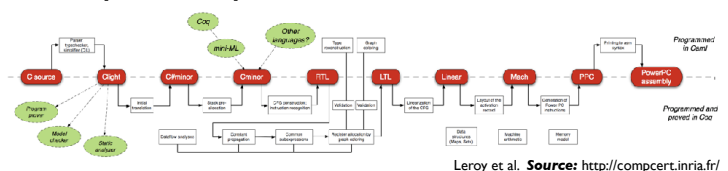
Compiler:



## State of the art

# CompCert

## CompCert C compiler



Compiles C source code to assembly.

Has **good performance numbers**

Proved correct in Coq.

<http://compcert.inria.fr/>

# CakeML compiler

Compiles CakeML concrete syntax to machine code.

Proved correct in HOL4.


Has mostly good performance numbers (later lecture)

Known as the first verified compiler to be bootstrapped.

I'm one of the six developers behind version 2 (diagram to the right).

<https://cakeml.org/>

larger on website



## A spectrum

robust, inflexible

proved to always  
work correctly

## Verified compilers

more flexible,  
but can be fragile

produces a proof for each run

## Proof-producing compilers

**Pilsner**

CompCert C compiler

## CakeML compiler

CompCertTSO

*Fiat*

**Cogent**

### Translation validation for a verified OS kernel

## Compiler verification summary

### Ingredients:

- a **formal logic** for the proofs
- **accurate models** of
  - the **source** language
  - the **target** language
  - the **compiler** algorithm

### Tools:

- a **proof assistant** (software)

**Method:**

- (interactively) prove a simulation relation

**Questions?** — contact me regarding MSc projects on this topic