

Towards **verified theorem provers**

Part I: proving the soundness of Milawa

QED+20, Vienna

Magnus O. Myreen — Computer Laboratory, University of Cambridge, UK

Jared Davis — Centaur Technology, Inc., Austin TX, USA

(also mentions work with Rob Arthan, Ramana Kumar, Michael Norrish, Scott Owens)

Verified theorem provers

Part 1: proving the soundness of Milawa

Part 2: (by Ramana) towards a verified impl. of HOL light

Part I: proving the soundness of Milawa



A self-verifying theorem prover

Jared Davis — PhD work

Jitawa
verified
LISP

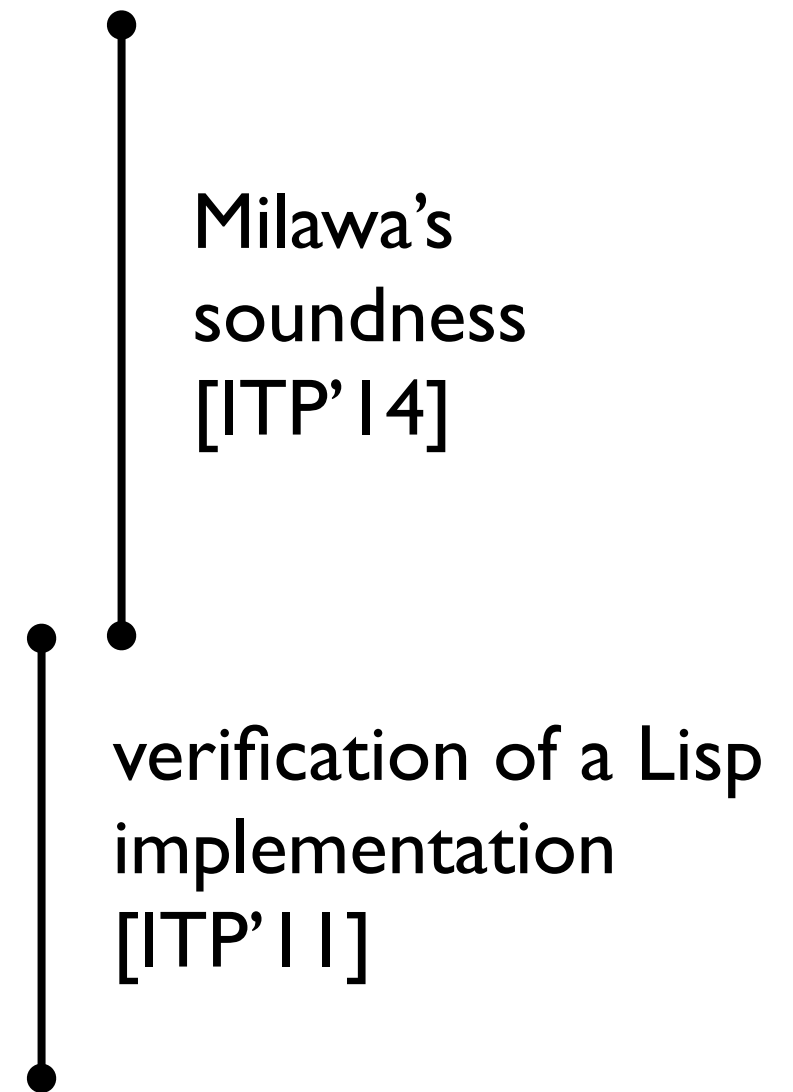
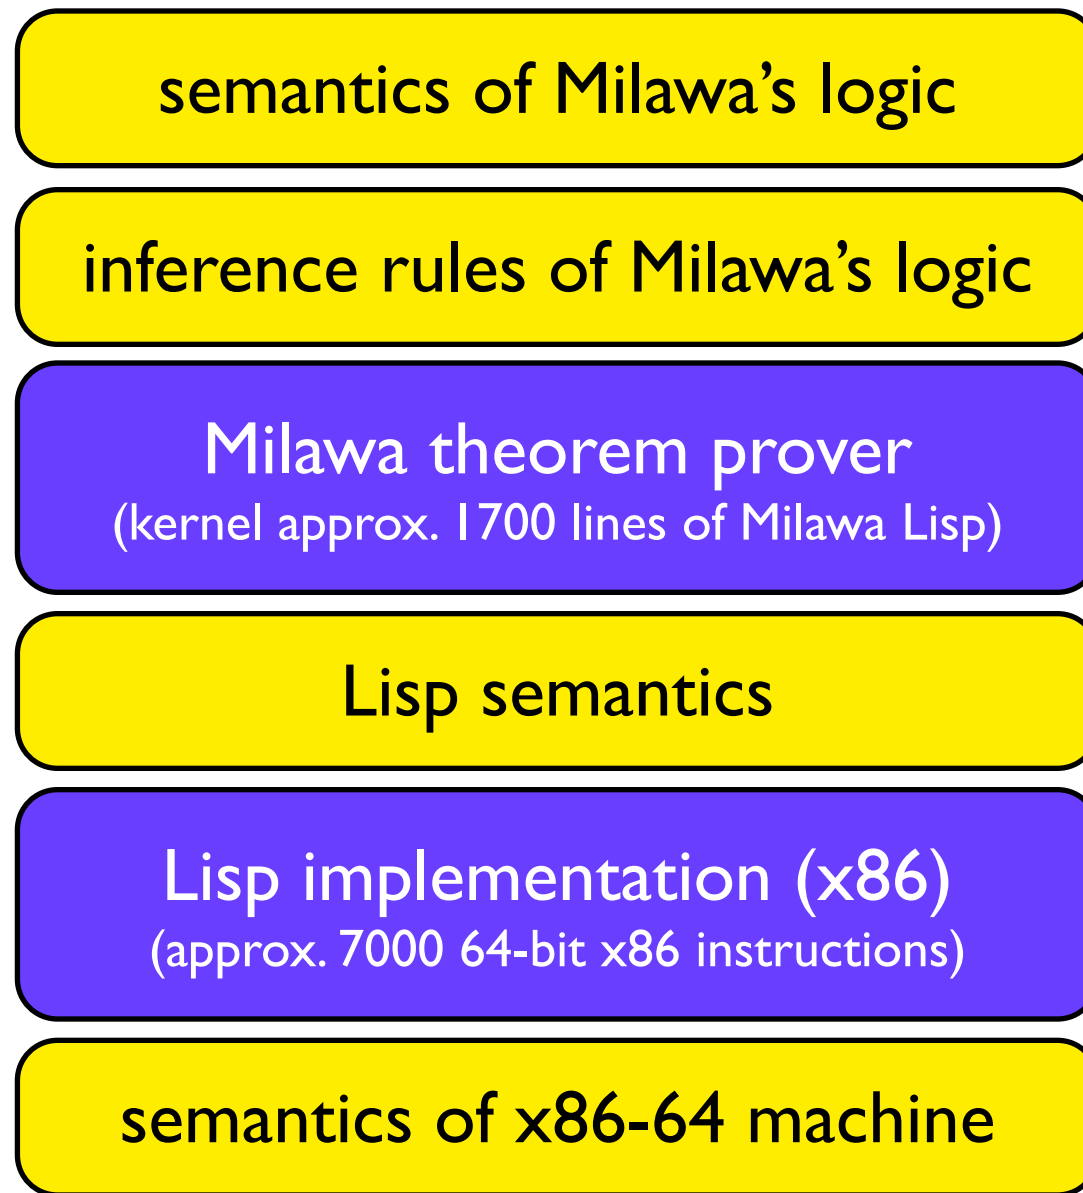
A verified runtime for a verified theorem prover

Magnus Myreen, Jared Davis — ITP'11

Proving Milawa sound



Jitawa
verified
LISP

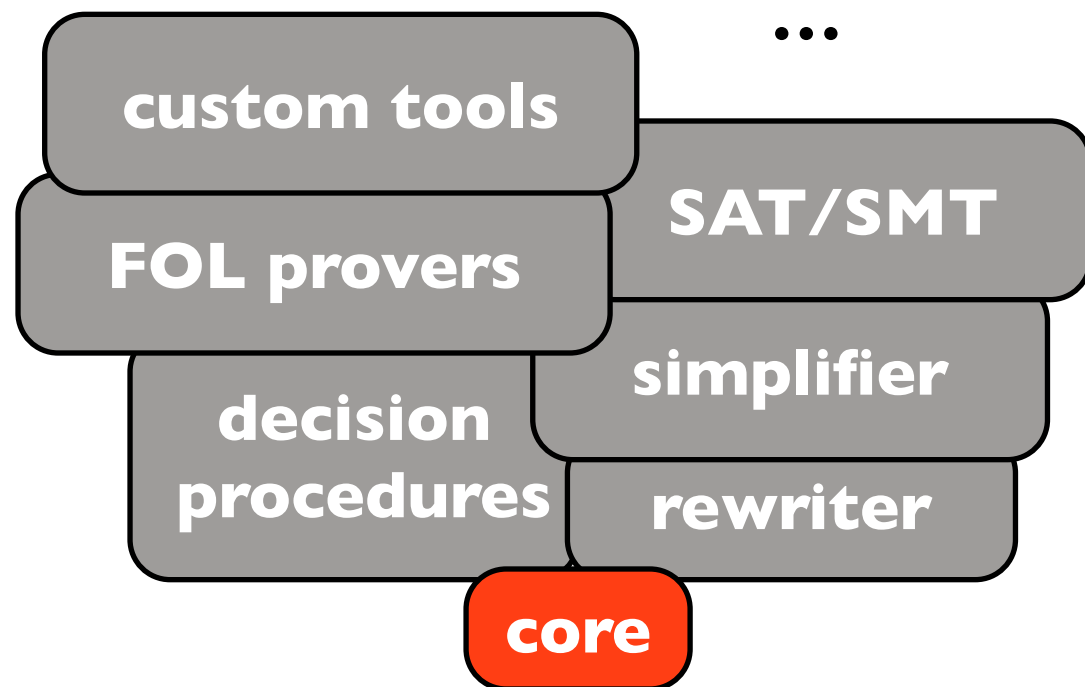


A very short introduction



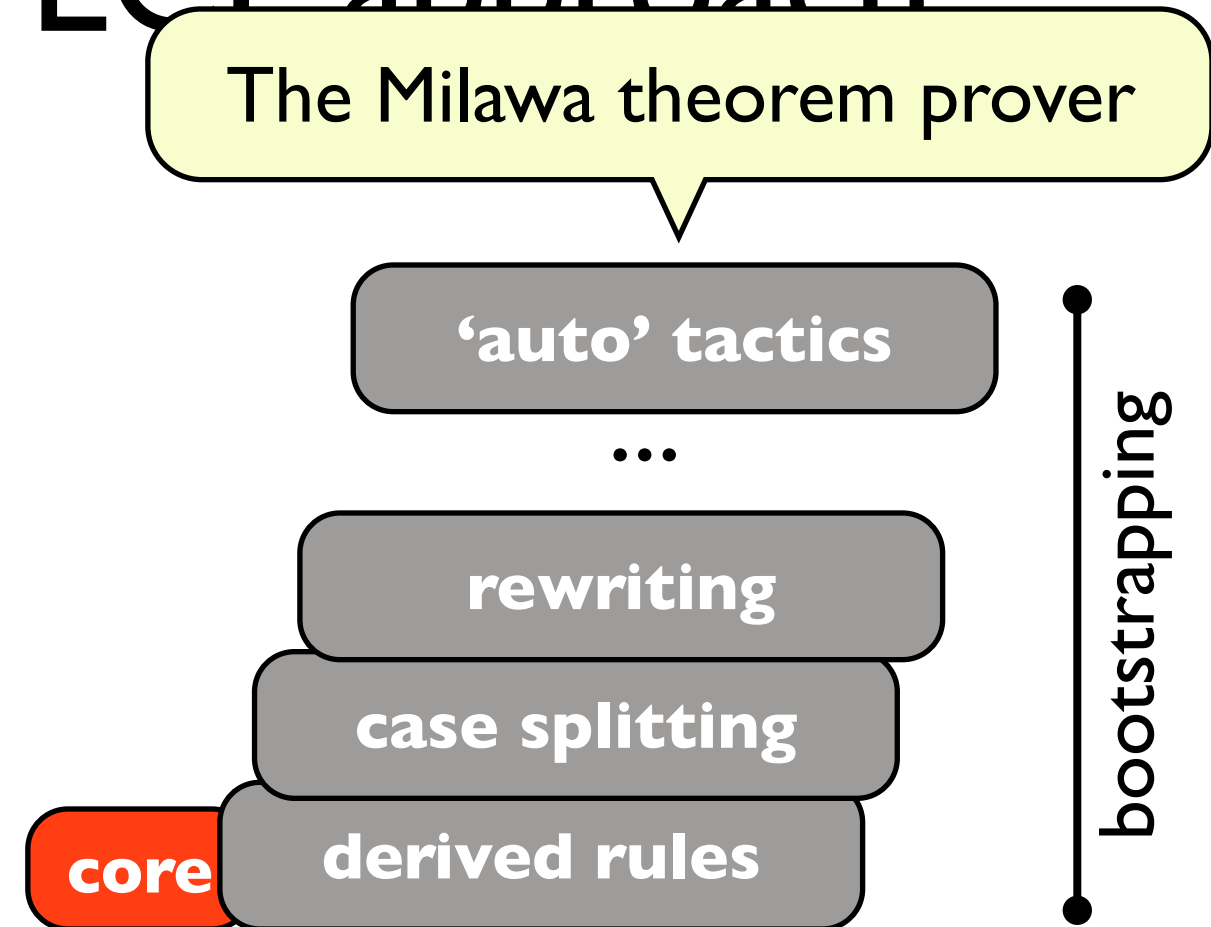
- Milawa is styled after theorem provers such as **NQTHM** and **ACL2**,
- has a **small trusted logical kernel** similar to LCF-style provers,
- ... but does **not suffer** the **performance hit of LCF's** fully expansive approach.

Comparison with LCF approach



LCF-style approach

- all proofs pass through the core's primitive inferences
- extensions steer the core



the Milawa approach

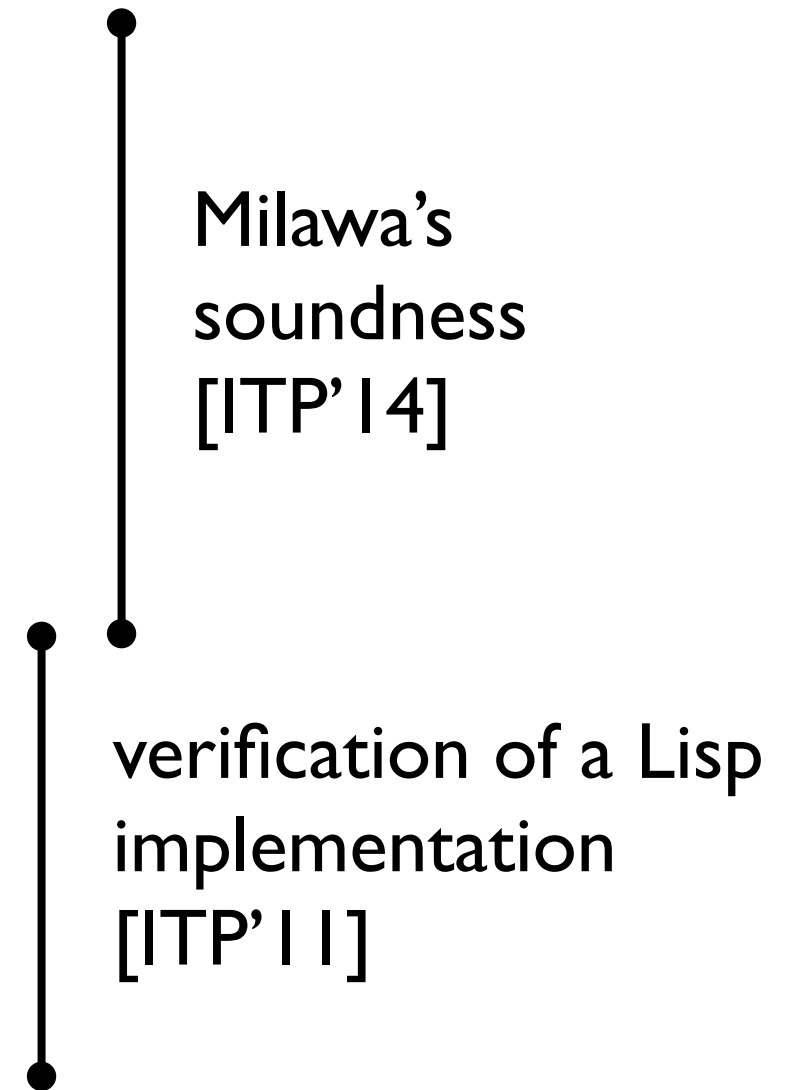
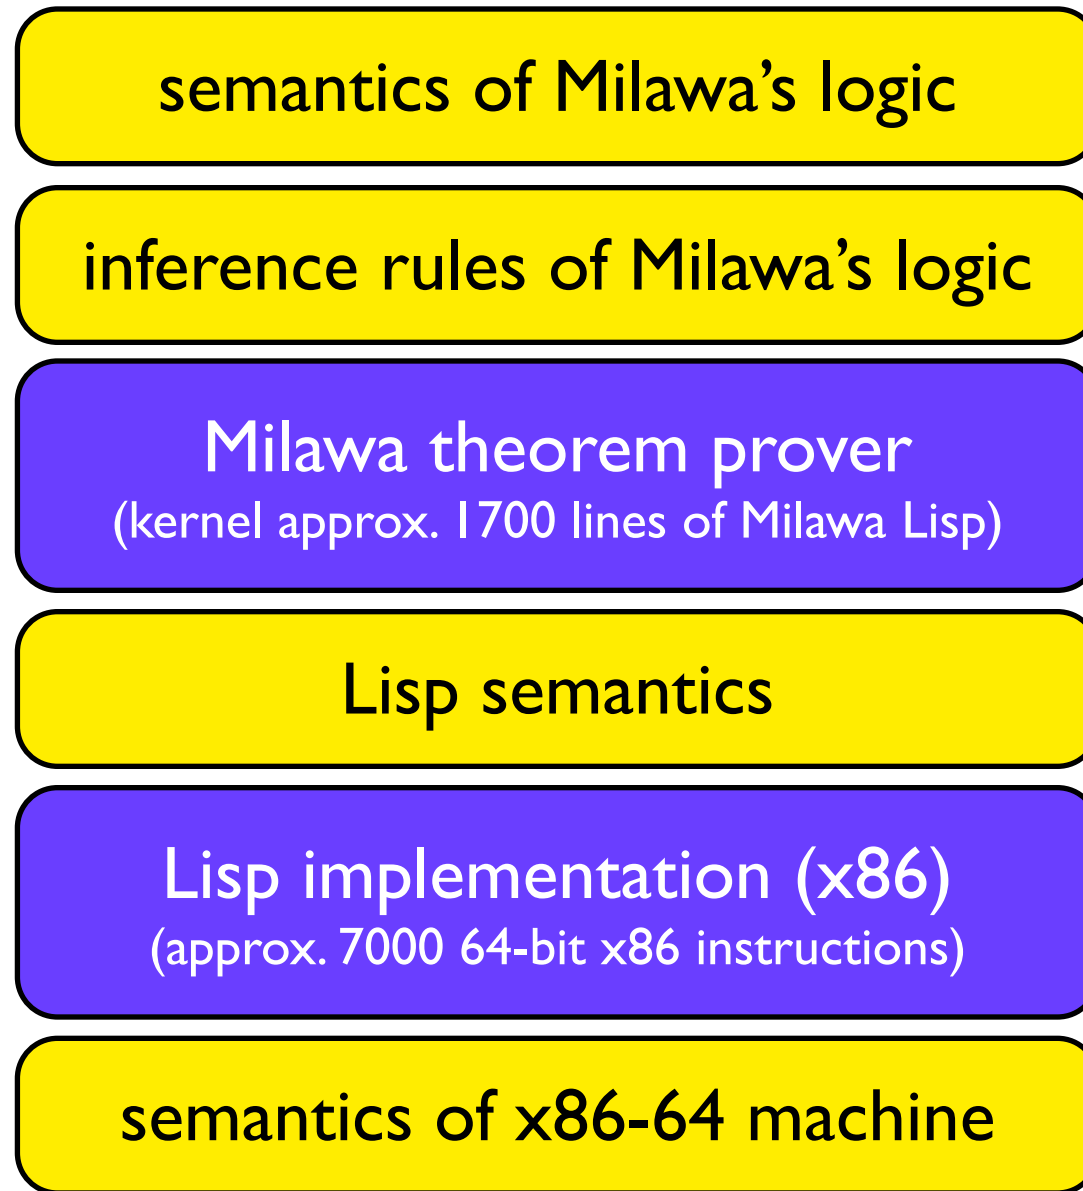
- all proofs must pass the core
- the **core proof checker** can be **replaced** at runtime

Brief Milawa demo

Proving Milawa sound



Jitawa
verified
LISP



Steps

A. formalise Milawa's logic

- ▶ syntax, semantics, inference, soundness

B. prove that Milawa's kernel is faithful to the logic

- ▶ run the Lisp parser (in the logic) on Milawa's kernel
- ▶ translate (with proof) deep embedding into shallow
- ▶ prove that Milawa's (reflective) kernel is faithful to logic

C. connect the verified Lisp implementation

- ▶ compose with the correctness thm from ITP'11

A—C combine to a top-level theorem that **relates** the **logic's semantics** with the execution of the **x86 machine code**.

This talk

A. formalise Milawa's logic

B. prove that Milawa's kernel is faithful to the logic

C. connect the verified Lisp implementation

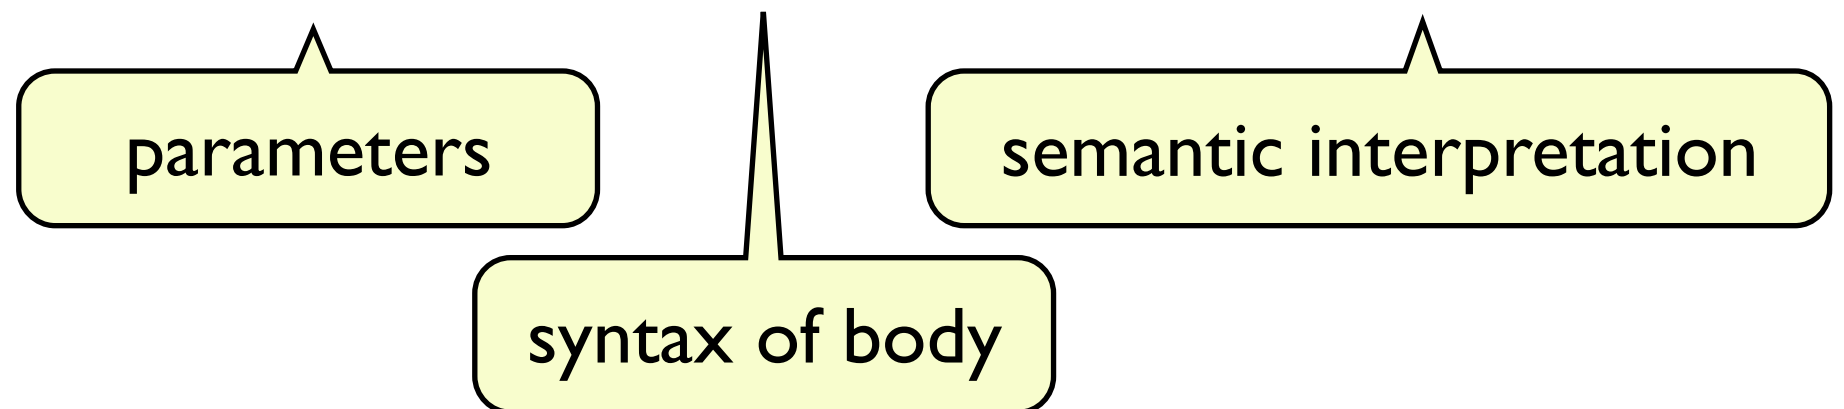
Syntax

$sexp ::= Val\ num \mid Sym\ string \mid Dot\ sexp\ sexp$	S-expression
$prim ::= If \mid Equal \mid Not \mid Symbolp \mid Symbol_less$ $\mid Natp \mid Add \mid Sub \mid Less \mid Consp \mid Cons$ $\mid Car \mid Cdr \mid Rank \mid Ord_less \mid Ordp$	
$func ::= PrimitiveFun\ prim$ $\mid Fun\ string$	primitive functions user-defined
$term ::= Const\ sexp$ $\mid Var\ string$ $\mid App\ func\ (term\ list)$ $\mid LamApp\ (string\ list)\ term\ (term\ list)$	constant S-expression variable function application λ formals body actuals
$formula ::= \neg formula$ $\mid formula \vee formula$ $\mid term = term$	negation disjunction term equality

Context

Syntax, semantics and inference rules depend on a context.

A context is a finite partial map
from *string* to *string list* \times *func_body* \times (*sexp list* \rightarrow *sexp*)



func_body ::= Body term
| Witness term string
| None

concrete term (e.g. recursive function)
property, element name
no function body given

Semantics

$$(\models_{\pi} p) = \text{formula_ok}_{\pi} p \wedge \forall i. \text{eval_formula } i \pi p$$

syntax makes sense

truth value

$$\text{eval_formula } i \pi (\neg p) = \neg(\text{eval_formula } i \pi p)$$

$$\text{eval_formula } i \pi (p \vee q) = \text{eval_formula } i \pi p \vee \text{eval_formula } i \pi q$$

$$\text{eval_formula } i \pi (x = y) = (\text{eval_term } i \pi x = \text{eval_term } i \pi y)$$

$$\text{eval_term } i \pi (\text{Const } c) = c$$

$$\text{eval_term } i \pi (\text{Var } v) = i(v)$$

$$\text{eval_term } i \pi (\text{App } f \ xs) = \text{eval_app } (f, \text{map } (\text{eval_term } i \pi) \ xs, \pi)$$

$$\text{eval_term } i \pi (\text{LambdaApp } vs \ x \ xs) = \text{let } ys = \text{map } (\text{eval_term } i \pi) \ xs \text{ in} \\ \text{eval_term } [vs \mapsto ys] \pi \ x$$

$$\text{eval_app } (\text{PrimitiveFun } p, \text{args}, \pi) = \text{eval_primitive } p \ \text{args}$$

$$\text{eval_app } (\text{Fun } name, \text{args}, \pi) = \text{let } (_, _, \text{interp}) = \pi(name) \text{ in} \\ \text{interp}(\text{args})$$

$$\text{eval_primitive } \text{Add } [\text{Val } 2, \text{Val } 3] = \text{Val } 5$$

$$\text{eval_primitive } \text{Add } [\text{Val } 2, \text{Sym } "a"] = \text{Val } 2$$

$$\text{eval_primitive } \text{Cons } [\text{Val } 2, \text{Sym } "a"] = \text{Dot } (\text{Val } 2) (\text{Sym } "a")$$

Well-formedness of context

Semantics only makes sense for well-formed contexts.

For every entry,

$$\pi(\textit{name}) = (\textit{formals}, \text{Body } \textit{body}, \textit{interp})$$

it must be that:

- ▶ the **formals** are all distinct
- ▶ the **body** is well-formed w.r.t. the **context**
- ▶ the **interpretation** satisfies the defining equation:

$$\forall i. \textit{interp}(\text{map } i \textit{formals}) = \text{eval_term } i \pi \textit{body}$$

Similarly for the other function types.

(a few of the) Inference rules

$$\frac{\vdash_{\pi} a \vee (b \vee c)}{\vdash_{\pi} (a \vee b) \vee c} \text{ (associativity)}$$

$$\frac{a \in \text{milawa_axioms}}{\vdash_{\pi} a} \text{ (basic axiom)}$$

facts about Lisp primitives

function definition in context

body of function

$$\frac{\pi(\text{name}) = (\text{formals}, \text{Body } \text{body}, \text{interp})}{\vdash_{\pi} \text{App } (\text{Fun } \text{name}) \text{ (map Var } \text{formals}) = \text{body}}$$

defining equation

Soundness of logic

Soundness of inference rules:

$$\forall \pi p. \text{context_ok } \pi \wedge (\vdash_{\pi} p) \implies (\models_{\pi} p)$$

- ▶ induction rule most interesting, Kaufmann&Slind [TPHOLs'07]

Soundness of definition mechanism:

$\forall \pi \text{ name formals body.}$

$\text{context_ok } \pi \wedge \text{definition_ok } (\text{name}, \text{formals}, \text{body}, \pi) \implies$

$\text{context_ok } (\pi[\text{name} \mapsto (\text{formals}, \text{body}, \text{new_interp } \pi \text{ name formals body})])$

- ▶ req. proving that termination conditions imply that a semantic interpretation exists as a function in HOL

This talk

A. formalise Milawa's logic

B. prove that Milawa's kernel is faithful to the logic

C. connect the verified Lisp implementation

This talk

A. formalise Milawa's logic

B. prove that Milawa's kernel is faithful to the logic

- ▶ run the Lisp parser (in the logic) on Milawa's kernel
- ▶ translate (with proof) deep embedding into shallow
- ▶ prove that Milawa's (reflective) kernel is faithful to logic

C. connect the verified Lisp implementation

Proving Milawa faithful to its logic

Verification must be w.r.t. semantics of Lisp [ITP'11].

Semantics of Lisp's read-eval-print loop:

1. **parse ASCII characters** into s-expressions
2. translate s-expressions into **program AST**
3. **evaluate** program **AST**
4. **print results**, goto 1.

Need to verify program down to **concrete source code**.

Steps towards an easier verification

- ▶ run the Lisp parser (in the logic) on Milawa's kernel

Each top-level function definition in ASCII

```
(defun lookup-safe (a x)
  (if (consp x)
      (if (equal a (car (car x)))
          (if (consp (car x))
              (car x)
              (cons (car (car x)) (cdr (car x))))
          (lookup-safe a (cdr x)))
      nil))
```

becomes a program AST

```
App Define [Const (Sym "LOOKUP-SAFE"), Const (...), Const (...)]
```

Steps towards an easier verification

When

App Define [Const (Sym "LOOKUP-SAFE"), Const (...), Const (...)]

is evaluated, the op. sem. adds a definition to its context:

function name: "LOOKUP-SAFE"
parameter list: "A", "X"
function body: If (App (PrimitiveFun Consp) [Var "X"])
 (If (App (PrimitiveFun Equal) [...])
 (If (App (PrimitiveFun Consp) [...] (...) (...))
 (App (Fun "LOOKUP-SAFE") [...]))
 (Const (Sym "NIL"))

We could do verification over this deep embedding.

...but a shallow embedding is easier to work with.

Steps towards an easier verification

function name: "LOOKUP-SAFE"
parameter list: "A", "X"
function body: If (App (PrimitiveFun Consp) [Var "X"]
 (If (App (PrimitiveFun Equal) [...])
 (If (App (PrimitiveFun Consp) [...] (...) (...))
 (App (Fun "LOOKUP-SAFE") [...])))
 (Const (Sym "NIL"))

We translate deep embedding into convenient shallow emb.
[ITP'12]

lookup_safe a x = if consp x then
 if $a = \text{car } (x)$ then
 if consp (x) then
 x
 else cons $(\text{car } (x))$ $(\text{cdr } (x))$
 else lookup_safe a $(\text{cdr } x)$
 else Sym "NIL"

Steps towards an easier verification

We translate deep embedding into convenient shallow emb.
[ITP'12]

```
lookup_safe a x = if consp x then
                  if a = car (car x) then
                    if consp (car x) then
                      car x
                    else cons (car (car x)) (cdr (car x))
                  else lookup_safe a (cdr x)
                  else Sym "NIL"
```

and produce a **certificate theorem** relating the deep and shallow embeddings.

... \implies (Fun "LOOKUP-SAFE", $[a, x]$, $state$) \xrightarrow{ap} (lookup_safe a x, $state$)

name in deep embedding

shallow embedding

Lisp semantics

Verification proof

- ▶ prove that Milawa's (reflective) kernel is faithful to logic

A routine verification exercise.

Points of interest:

Milawa's **initial proof checker** was a large function

Top-level loop has complicated invariant, relates:

- ▶ program state
- ▶ current Lisp op.sem. state
- ▶ logical context

Bugs found? Yes, two very minor (no soundness bugs)

Verification proof

Theorem:

$\exists ans\ k\ output\ ok.$

$milawa_main\ cmds\ init_state = (ans, (k, output, ok)) \wedge$

$(ok \implies (ans = \text{Sym "SUCCESS"})) \wedge$

$\text{let } result = \text{compute_output } cmds \text{ in}$

$\text{every_line } line_ok\ result \wedge$

$output = \text{output_string } result)$

where

$line_ok\ (\pi, l) = (l = \text{"NIL"}) \vee$
 $(\exists n. (l = \text{"(PRINT (n ...))"}) \wedge \text{is_number } n) \vee$
 $(\exists \phi. (l = \text{"(PRINT (THEOREM } \phi))"}) \wedge \text{context_ok } \pi \wedge \models_{\pi} \phi)$

This talk

A. formalise Milawa's logic

B. prove that Milawa's kernel is faithful to the logic

C. connect the verified Lisp implementation

Correctness of Jitawa Lisp [ITP'11]

Top-level correctness theorem:

$\{ \text{init_state } input * \text{pc } pc * \langle \text{terminates_for } input \rangle \}$

$pc : \text{code_for_entire_jitawa_implementation}$

$\{ \text{error_message } \vee \exists output. \langle ([], input) \xrightarrow{\text{exec}} (output, \text{true}) \rangle * \text{final_state } output \}$

Correctness of Jitawa Lisp [ITP'11]

There must be enough memory and I/O assumptions must hold.

This machine-code Hoare triple holds only for terminating executions.

$\{ \text{init_state } input * pc \text{ } pc * \langle \text{terminates_for } input \rangle \}$

$pc : \text{code_for_entire_jitawa_implementation}$

list of numbers

$\{ \text{error_message } \vee \exists output. \langle ([], input) \xrightarrow{\text{exec}} (output, \text{true}) \rangle * \text{final_state } output \}$

Each execution is allowed to fail with an error message.

If there is no error message, then the result is described by the high-level op. semantics.

Theorem: Milawa is sound down to x86

There must be enough memory and input is Milawa's kernel followed by call to main for some *input*.

\forall *input pc*.

```
{ init_state (milawa_implementation ++ "(milawa-main 'input)") * pc pc }  
pc : code_for_entire_jitawa_implementation  
{ error_message  $\vee$  (let result = compute_output (parse input) in  
  <every_line line_ok result> *  
  final_state (output_string result ++ "SUCCESS")) }
```

Machine code terminates either with error message, or ...

... output lines that are all true w.r.t. the semantics of the logic.

Running Milawa on Jitawa

Running Milawa's 4-gigabyte bootstrap process:

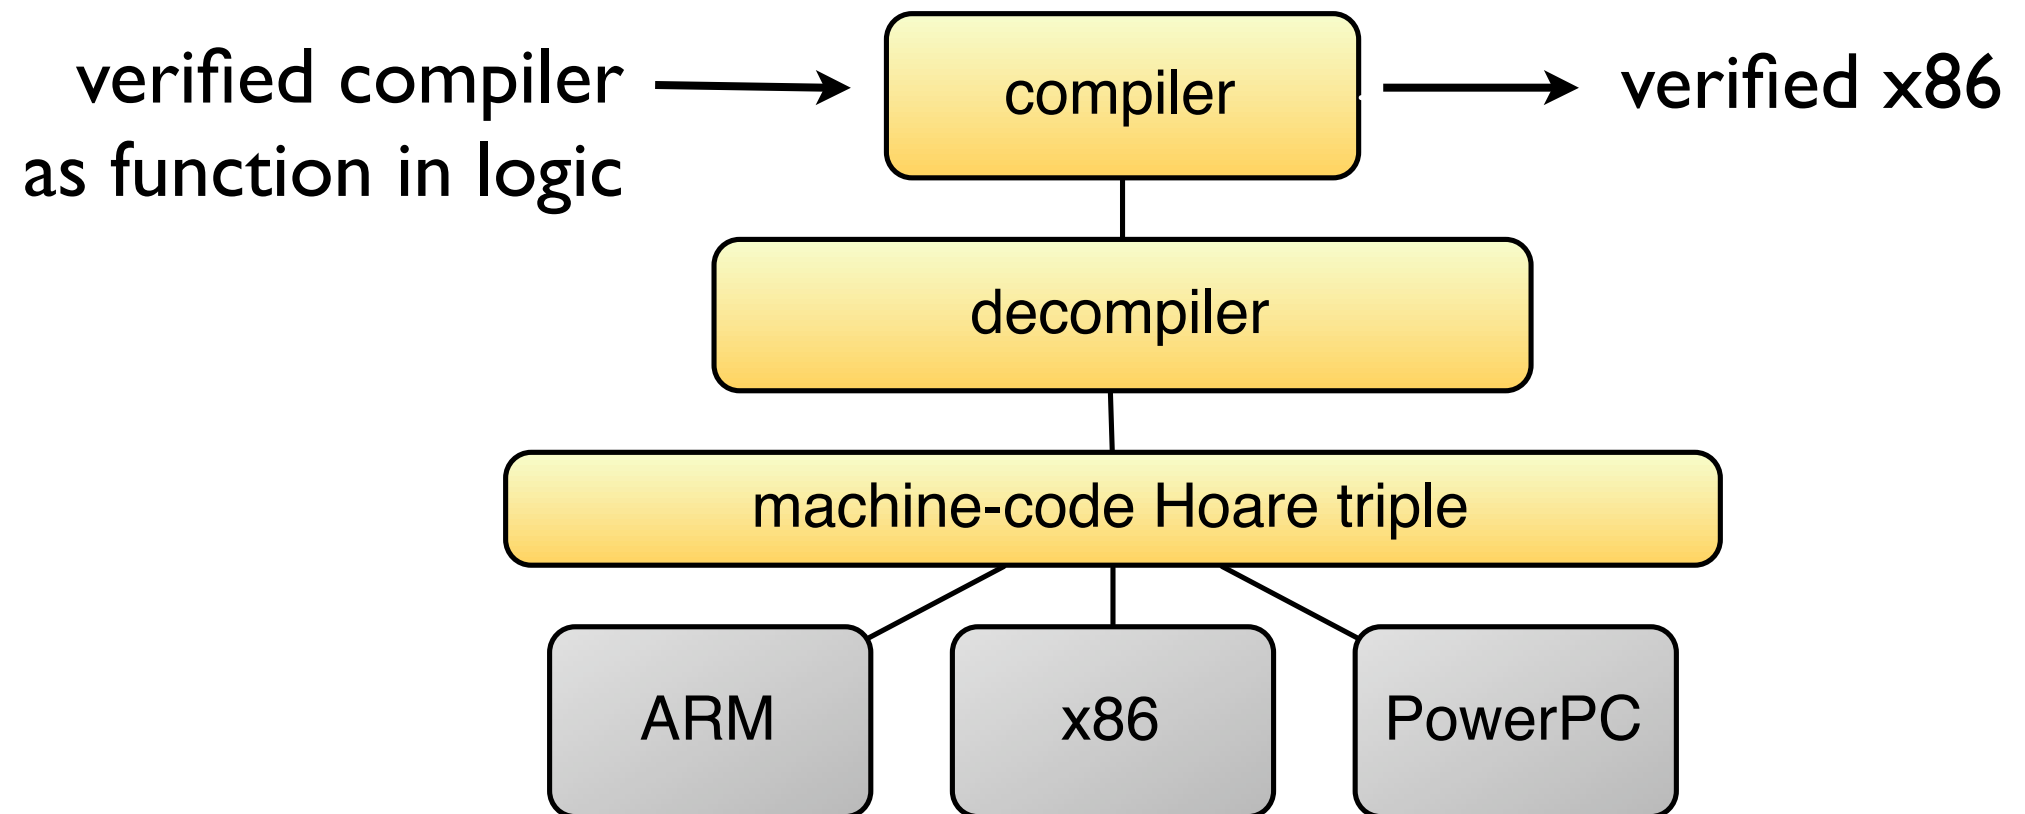
CCL	16 hours	Jitawa's compiler performs almost no optimisations.
SBCL	22 hours	
Jitawa	128 hours (8x slower than CCL)	

Parsing the 4 gigabyte input:

CCL	716 seconds (9x slower than Jitawa)
Jitawa	79 seconds

Looking back...

The x86 for the compile function was produced as follows:



Very cumbersome....

...should have compiled the verified compiler using itself!

Bootstrapping the compiler

Instead: we bootstrap the verified compile function, we evaluate the compiler on a deep embedding of itself within the logic:

EVAL ``compile COMPILE``

derives a theorem:

in Lisp (eval '(compile compile)) ?

compile COMPILE = **compiler-as-machine-code**

The first(?) bootstrapping of a formally verified compiler.



Ramana Kumar
(Uni. Cambridge)



Magnus Myreen
(Uni. Cambridge)



Michael Norrish
(NICTA, ANU)



Scott Owens
(Uni. Kent)

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 machine code. Our correctness theorem ensures that these results permitted

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 compilers for general-purpose languages has addressed all the following dimensions: one, the compilation

Summary

The top-level theorem:

relates the **logic's semantics**
with the execution of the **x86 machine code**.

Questions?

Steps:

A. formalise Milawa's logic

- ▶ syntax, semantics, inference, soundness

B. prove that Milawa's kernel is faithful to the logic

- ▶ run the Lisp parser (in the logic) on Milawa's kernel
- ▶ translate (with proof) deep embedding into shallow
- ▶ prove that Milawa's (reflective) kernel is faithful to logic

C. connect the verified Lisp implementation

- ▶ compose with the correctness thm from ITP'11