

A verified runtime for a verified theorem prover

Magnus Myreen¹ and Jared Davis²

¹ University of Cambridge, UK

² Centaur Technology, USA

ITP 2011

Two projects meet

Jared Davis

A self-verifying
theorem prover



Milawa

Magnus Myreen

Verified Lisp
implementations

verified **LISP** on
ARM, x86, PowerPC

Two projects meet

My theorem prover is written in Lisp.
Can I try your verified Lisp interpreter?

Jared Davis

A self-verifying
theorem prover



Milawa

Magnus Myreen

Verified Lisp
implementations

verified **LISP** on
ARM, x86, PowerPC

Two projects meet

My theorem prover is written in Lisp.
Can I try your verified Lisp interpreter?

Jared Davis

A self-verifying
theorem prover



Milawa

Sure, try it.

Magnus Myreen

Verified Lisp
implementations

verified **LISP** on
ARM, x86, PowerPC

Two projects meet

My theorem prover is written in Lisp.
Can I try your verified Lisp interpreter?

Sure, try it.

Does your Lisp support ..., ... and ...?

Jared Davis

Magnus Myreen

A self-verifying
theorem prover

Verified Lisp
implementations



Milawa

verified **LISP** on
ARM, x86, PowerPC

Two projects meet

My theorem prover is written in Lisp.
Can I try your verified Lisp interpreter?

Sure, try it.

Does your Lisp support ..., ... and ...?

No, but it could ...

Jared Davis

Magnus Myreen

A self-verifying
theorem prover

Verified Lisp
implementations



Milawa

verified **LISP** on
ARM, x86, PowerPC

Running Milawa



Milawa



verified **LISP** on
ARM, x86, PowerPC
(TPHOLs 2009)

Running Milawa

Milawa's bootstrap proof:



verified **LISP** on
ARM, x86, PowerPC
(TPHOLs 2009)

Running Milawa



verified **LISP** on
ARM, x86, PowerPC
(TPHOLs 2009)

Milawa's bootstrap proof:

- ▶ 4 gigabyte proof file:
>500 million unique conseqs

Running Milawa



verified **LISP** on
ARM, x86, PowerPC
(TPHOLs 2009)

Milawa's bootstrap proof:

- ▶ 4 gigabyte proof file:
>500 million unique conseqs
- ▶ takes 16 hours to run on a
state-of-the-art runtime (CCL)

Running Milawa



verified **LISP** on
ARM, x86, PowerPC
(TPHOLs 2009)

Milawa's bootstrap proof:

- ▶ 4 gigabyte proof file:
>500 million unique conseqs
- ▶ takes 16 hours to run on a
state-of-the-art runtime (CCL)

← hopelessly “toy”

Running Milawa



Jitawa: verified **LISP**
with JIT compiler

Milawa's bootstrap proof:

- ▶ 4 gigabyte proof file:
>500 million unique conseqs
- ▶ takes 16 hours to run on a state-of-the-art runtime (CCL)

Contribution:

- ▶ a new verified Lisp which is able to host the Milawa thm prover

Outline

Part 1: **Milawa**

Part 2: **Its new verified runtime**

Part 3: **Mini-demos, measurements**

A short introduction to



- 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.

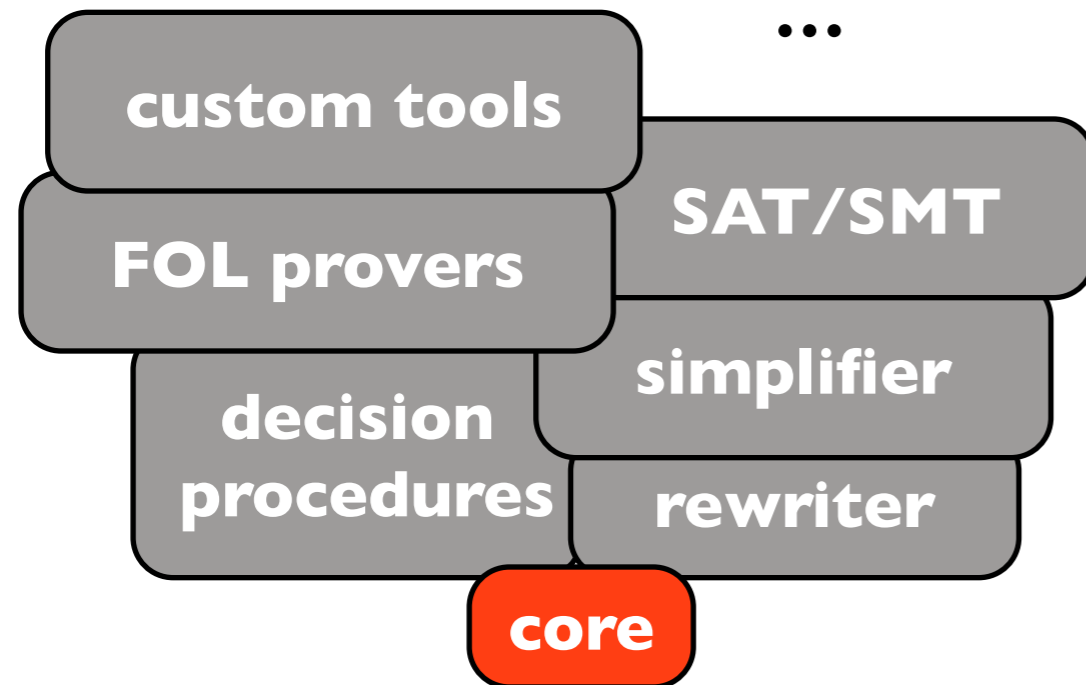
Comparison with LCF approach

core

LCF-style approach

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

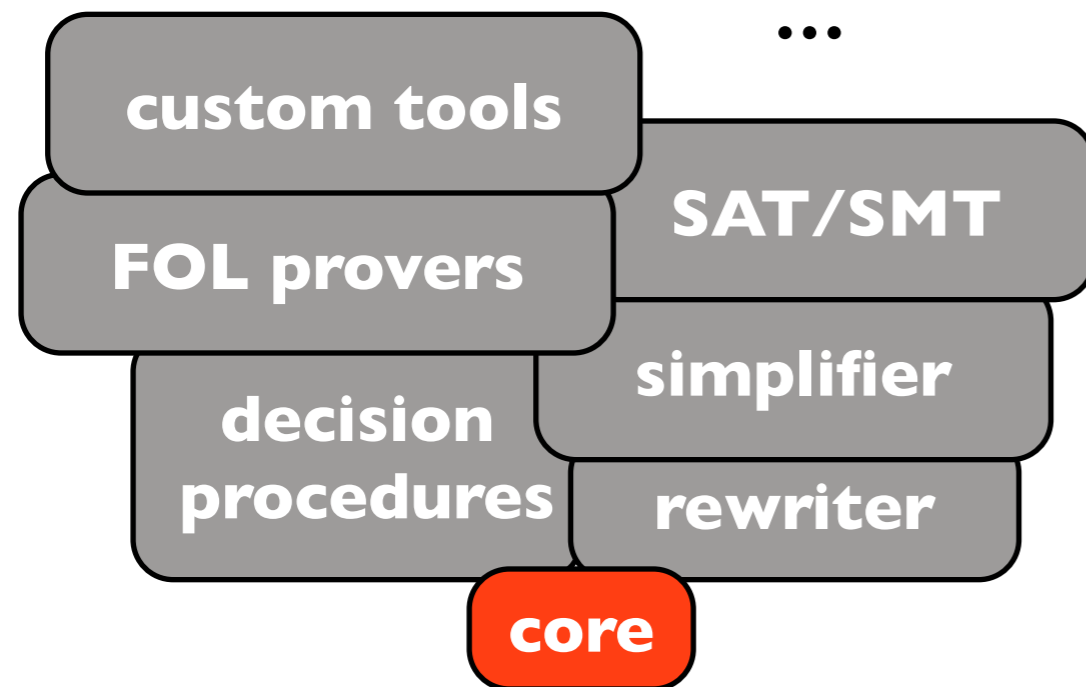
Comparison with LCF approach



LCF-style approach

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

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 can be reflectively extended at runtime

Comparison with LCF approach



LCF-style approach


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



the Milawa approach

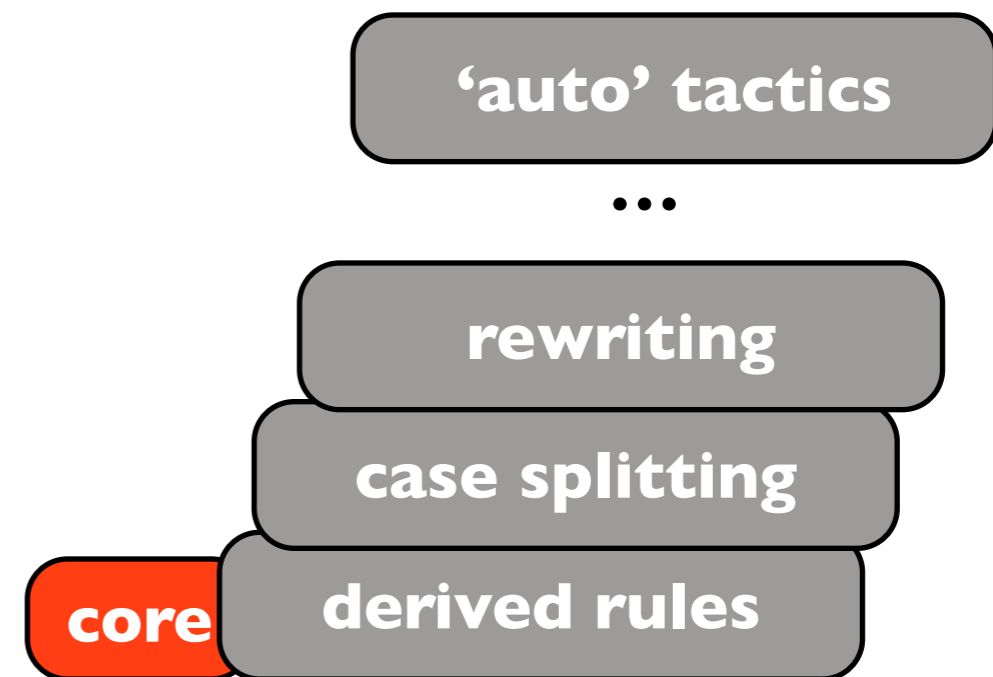
- all proofs must pass the core
- the core can be reflectively extended at runtime

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 can be reflectively extended at runtime

Bootstrapping Milawa

Output from Milawa's bootstrap proof:

```
(PRINT (1 VERIFY THEOREM-SUBSTITUTE-INTO-NOT-PEQUAL))
(PRINT (2 VERIFY THEOREM-NOT-T-OR-NOT-NIL))
(PRINT (3 DEFINE NOT))
(PRINT (4 VERIFY NOT))
(PRINT (5 DEFINE IFF))
(PRINT (6 VERIFY IFF))
(PRINT (7 VERIFY THEOREM-COMMUTATIVITY-OF-PEQUAL))
...
(PRINT (4611 VERIFY |INSTALL-NEW-PROOFP-LEVEL2.PROOFP|))
(PRINT (4612 SWITCH |LEVEL2.PROOFP|))
(PRINT (4613 VERIFY |BUST-UP-LOGIC.FUNCTION-ARGS-EXPENSIVE|))
...
(PRINT (15685 VERIFY |INSTALL-NEW-PROOFP-LEVEL11.PROOFP|))
(PRINT (15686 SWITCH |LEVEL11.PROOFP|))
...
SUCCESS
```

Bootstrapping Milawa

Output from ~~Milawa's bootstrap proof~~

starts with very basic definitions and lemmas

```
(PRINT (1 VERIFY THEOREM-SUBSTITUTE-INTO-NOT-PEQUAL))
(PRINT (2 VERIFY THEOREM-NOT-T-OR-NOT-NIL))
(PRINT (3 DEFINE NOT))
(PRINT (4 VERIFY NOT))
(PRINT (5 DEFINE IFF))
(PRINT (6 VERIFY IFF))
(PRINT (7 VERIFY THEOREM-COMMUTATIVITY-OF-PEQUAL))
...
(PRINT (4611 VERIFY |INSTALL-NEW-PROOFP-LEVEL2.PROOFP|))
(PRINT (4612 SWITCH |LEVEL2.PROOFP|))
(PRINT (4613 VERIFY |BUST-UP-LOGIC.FUNCTION-ARGS-EXPENSIVE|))
...
(PRINT (15685 VERIFY |INSTALL-NEW-PROOFP-LEVEL11.PROOFP|))
(PRINT (15686 SWITCH |LEVEL11.PROOFP|))
...
SUCCESS
```

Bootstrapping Milawa

Output from Milawa's bootstrap proof:

starts with very basic definitions and lemmas

```
(PRINT (1 VERIFY THEOREM-SUBSTITUTE-INTO-NOT-PEQUAL))
```

```
(PRINT (2 VERIFY THEOREM-NOT-T-OR-NOT-NIL))
```

```
(PRINT (3 DEFINE NOT))
```

```
(PRINT (4 VERIFY NOT))
```

```
(PRINT (5 DEFINE IF))
```

```
(PRINT (6 VERIFY IF))
```

```
(PRINT (7 VERIFY THEOREM-COMMUTATIVITY-OF-PEQUAL))
```

```
...
```

```
(PRINT (4611 VERIFY |INSTALL-NEW-PROOFP-LEVEL2.PROOFP|))
```

```
(PRINT (4612 SWITCH |LEVEL2.PROOFP|))
```

```
(PRINT (4613 VERIFY |BUST-UP-LOGIC.FUNCTION-ARGS-EXPENSIVE|))
```

```
...
```

```
(PRINT (15685 VERIFY |INSTALL-NEW-PROOFP-LEVEL11.PROOFP|))
```

```
(PRINT (15686 SWITCH |LEVEL11.PROOFP|))
```

```
...
```

```
SUCCESS
```

up to this point the original core is used

Bootstrapping Milawa

Output from Milawa's bootstrap proof:

starts with very basic definitions and lemmas

```
(PRINT (1 VERIFY THEOREM-SUBSTITUTE-INTO-NOT-PEQUAL))
```

```
(PRINT (2 VERIFY THEOREM-NOT-T-OR-NOT-NIL))
```

```
(PRINT (3 DEFINE NOT))
```

```
(PRINT (4 VERIFY NOT))
```

```
(PRINT (5 DEFINE IF))
```

```
(PRINT (6 VERIFY IF))
```

```
(PRINT (7 VERIFY THEOREM-NOT-PEQUAL))
```

up to this point the original core is used

this event switches to a new extended core

```
...
```

```
(PRINT (4611 VERIFY |INSTALL-NEW-PROOFP-LEVEL2.PROOFP|))
```

```
(PRINT (4612 SWITCH |LEVEL2.PROOFP|))
```

```
(PRINT (4613 VERIFY |BUST-UP-LOGIC.FUNCTION-ARGS-EXPENSIVE|))
```

```
...
```

```
(PRINT (15685 VERIFY |INSTALL-NEW-PROOFP-LEVEL11.PROOFP|))
```

```
(PRINT (15686 SWITCH |LEVEL11.PROOFP|))
```

```
...
```

```
SUCCESS
```

Bootstrapping Milawa

Output from Milawa's bootstrap proof:

starts with very basic definitions and lemmas

```
(PRINT (1 VERIFY THEOREM-SUBSTITUTE-INTO-NOT-PEQUAL))
```

```
(PRINT (2 VERIFY THEOREM-NOT-T-OR-NOT-NIL))
```

```
(PRINT (3 DEFINE NOT))
```

```
(PRINT (4 VERIFY NOT))
```

```
(PRINT (5 DEFINE IF))
```

```
(PRINT (6 VERIFY IF))
```

```
(PRINT (7 VERIFY THEOREM-NOT-OR-NOT-NIL))
```

```
...
```

```
(PRINT (4611 VERIFY (INSTALL-NEW-PROOFP-LEVEL2.PROOFP)))
```

```
(PRINT (4612 SWITCH |LEVEL2.PROOFP|))
```

```
(PRINT (4613 VERIFY (BUST-UP-LOGIC.FUNCTION-ARGS-EXPENSIVE)))
```

```
...
```

```
(PRINT (15685 VERIFY (INSTALL-NEW-PROOFP-LEVEL11.PROOFP)))
```

```
(PRINT (15686 SWITCH |LEVEL11.PROOFP|))
```

```
...
```

```
SUCCESS
```

up to this point the original core is used

this event switches to a new extended core

the extended core is used from now onwards

Bootstrapping Milawa

Output from Milawa's bootstrap proof:

starts with very basic definitions and lemmas

```
(PRINT (1 VERIFY THEOREM-SUBSTITUTE-INTO-NOT-PEQUAL))  
(PRINT (2 VERIFY THEOREM-NOT-T-OR-NOT-NIL))  
(PRINT (3 DEFINE NOT))  
(PRINT (4 VERIFY NOT))  
(PRINT (5 DEFINE IF))  
(PRINT (6 VERIFY IF))  
(PRINT (7 VERIFY THEOREM-NOT-T-OR-NOT-NIL))
```

up to this point the original core is used

this event switches to a new extended core

```
...  
(PRINT (4611 VERIFY (INSTALL-NEW-PROOFP-LEVEL2.PROOFP I))  
(PRINT (4612 SWITCH (LEVEL2.PROOFP I))  
(PRINT (4613 VERIFY (BUST-UP-LOGIC.FUNCTION-ARGS-EXPENSIVE I))
```

the extended core is used from now onwards

```
...  
(PRINT (15685 VERIFY (INSTALL-NEW-PROOFP-LEVEL11.PROOFP I))  
(PRINT (15686 SWITCH (LEVEL11.PROOFP I))
```

10 core extensions during bootstrap

```
...  
SUCCESS
```

Milawa's core extensions

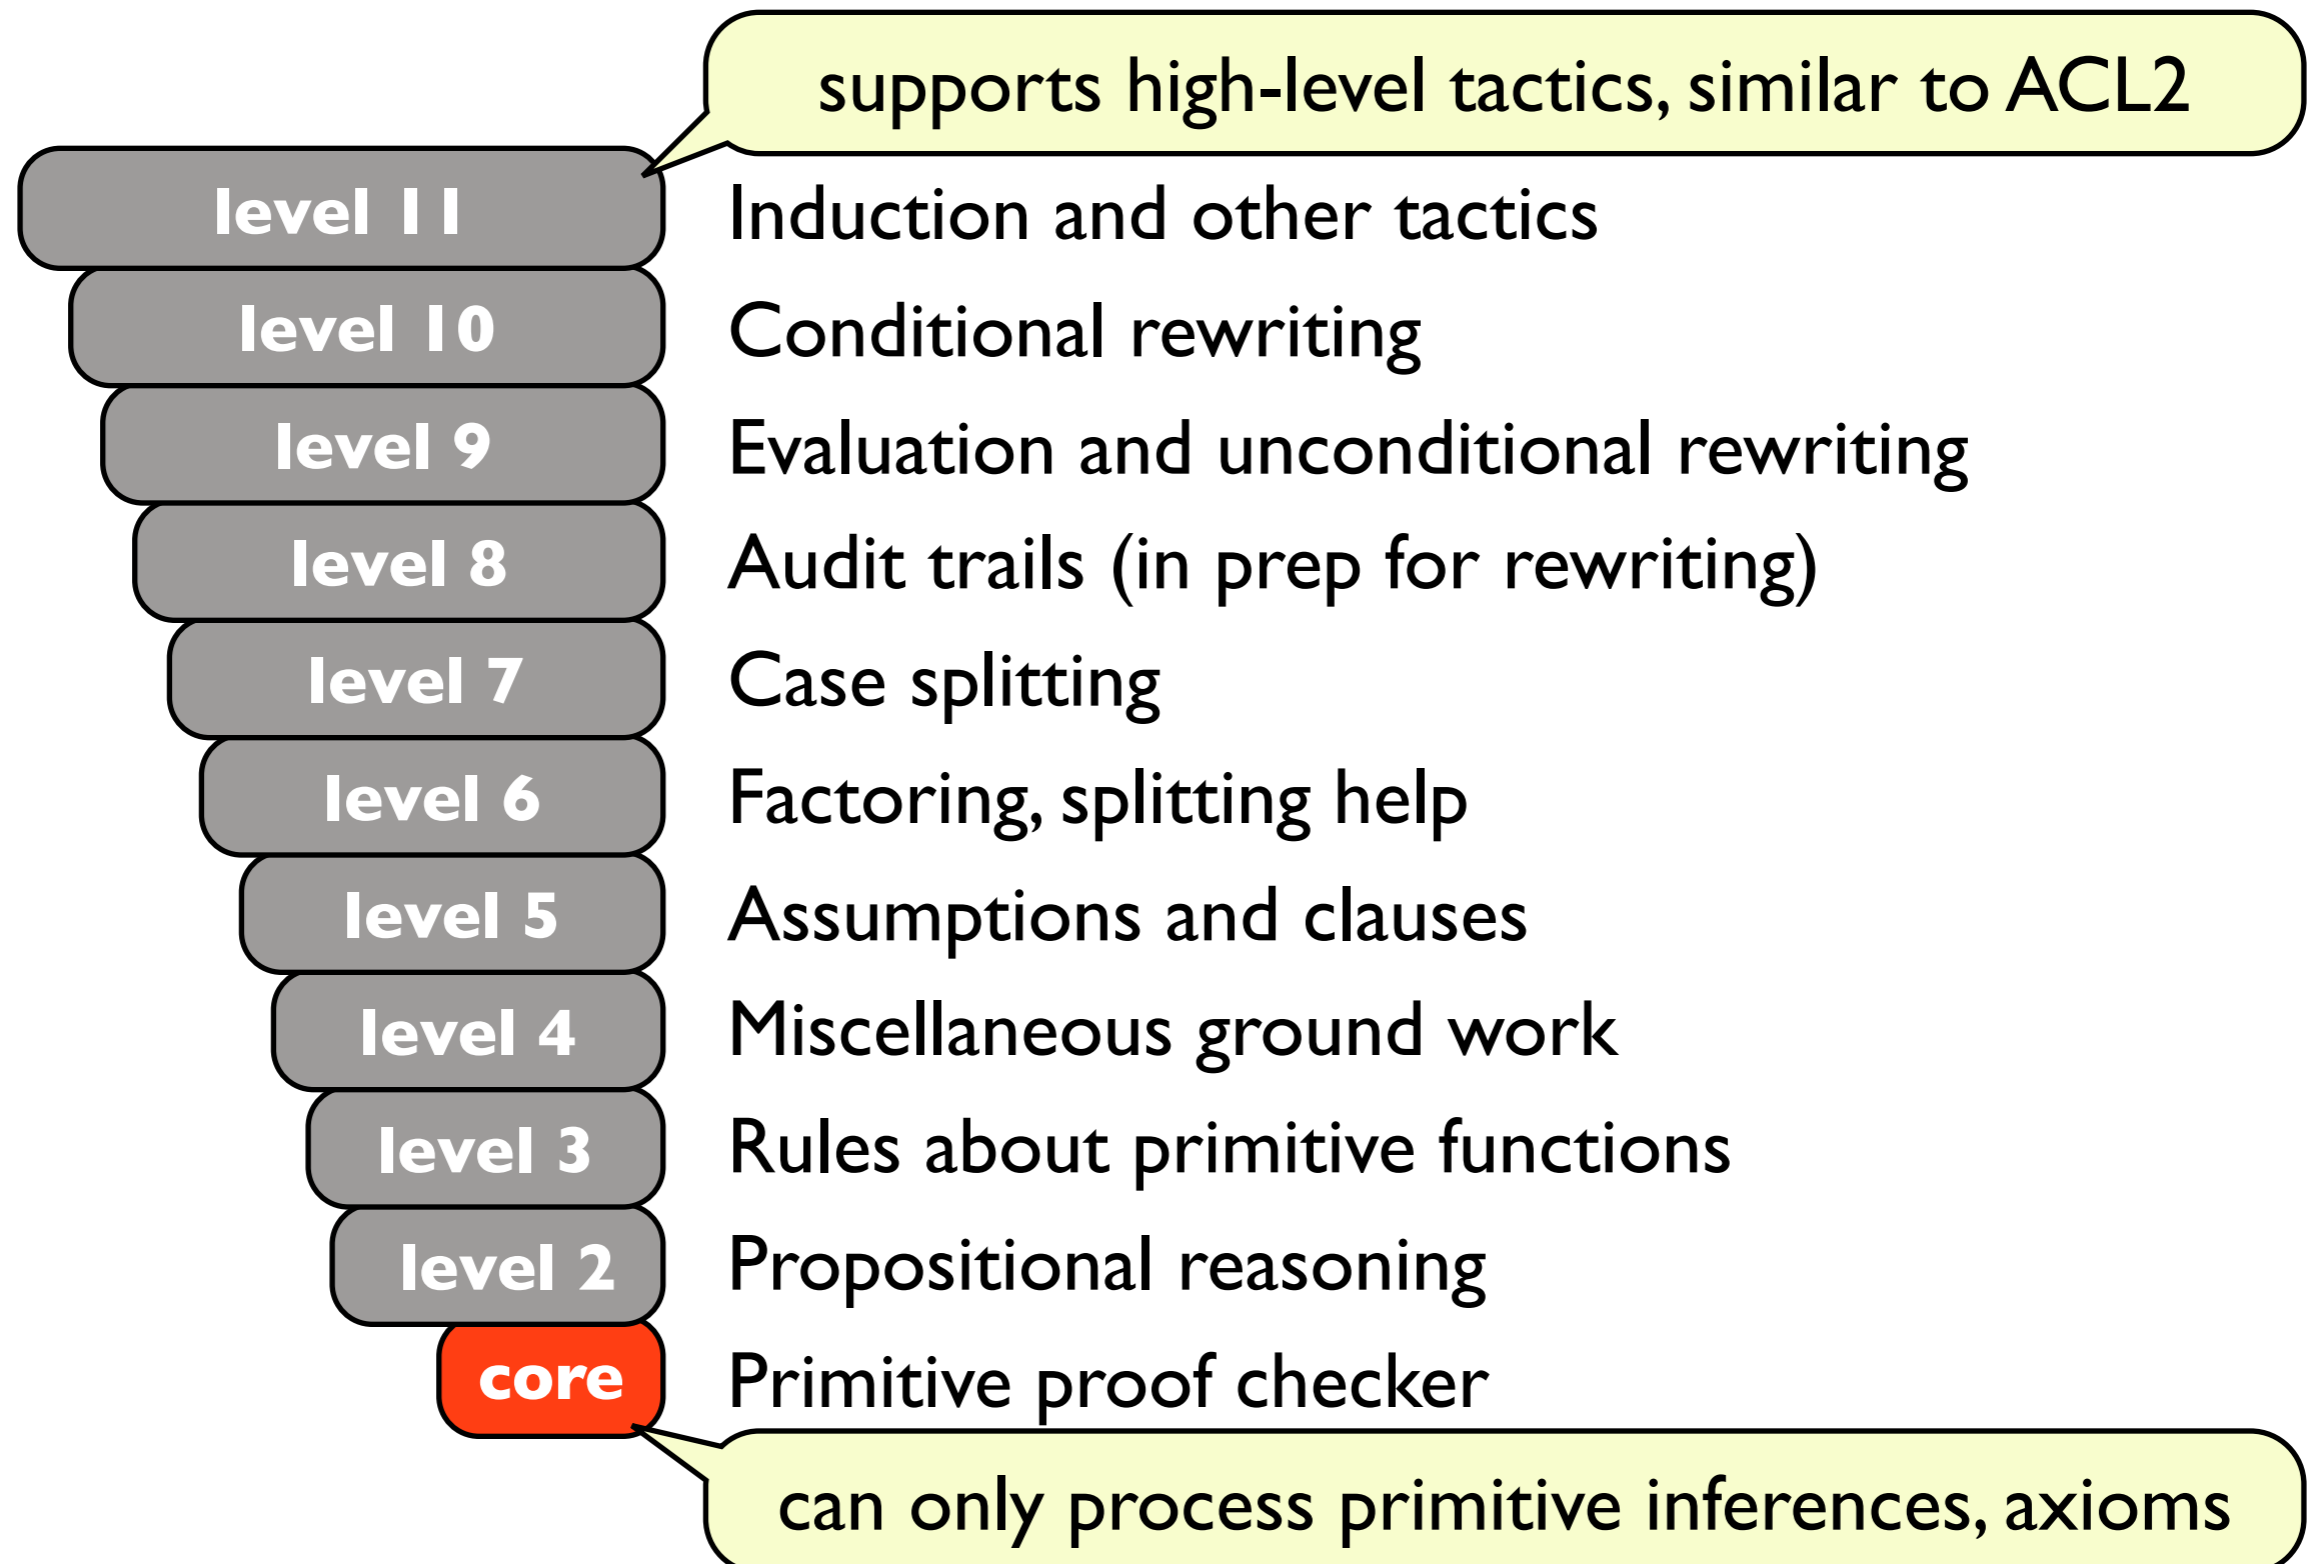
core

Milawa's core extensions

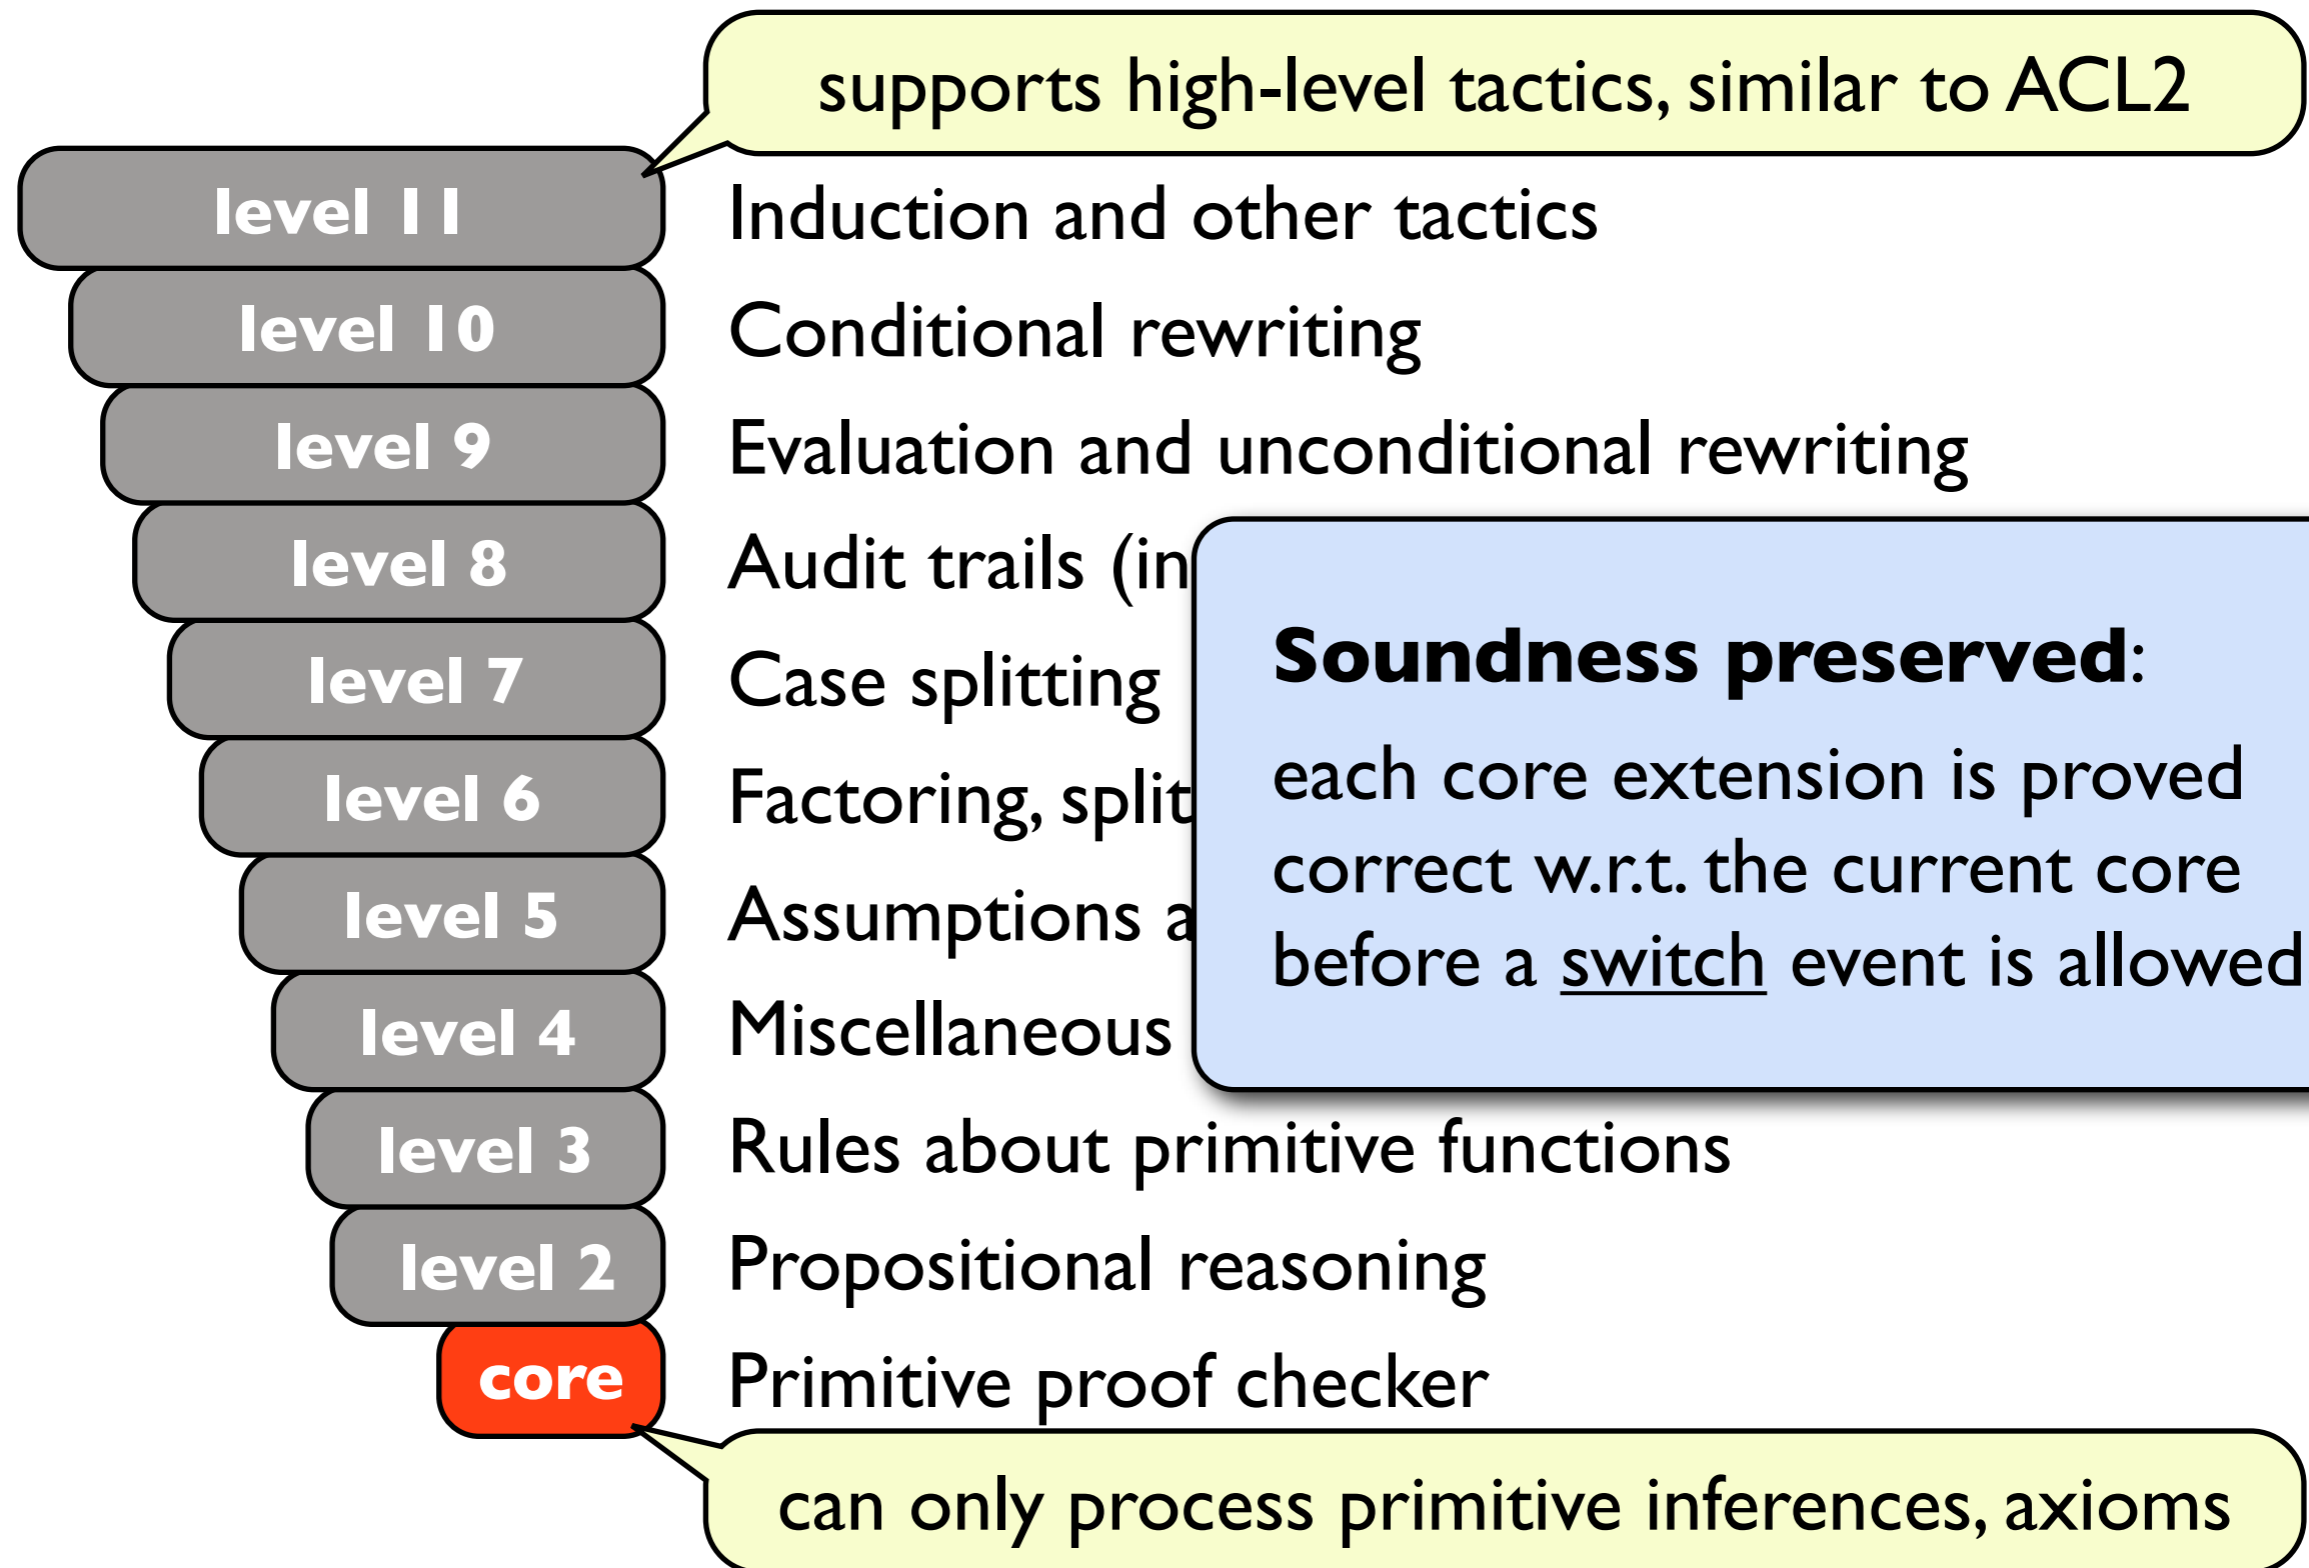
core

can only process primitive inferences, axioms

Milawa's core extensions



Milawa's core extensions



Milawa's logic

Prop. Schema

$$\frac{}{\neg A \vee A}$$

Contraction

$$\frac{A \vee A}{A}$$

Expansion

$$\frac{A}{B \vee A}$$

Associativity

$$\frac{A \vee (B \vee C)}{(A \vee B) \vee C}$$

Cut

$$\frac{A \vee B \quad \neg A \vee C}{B \vee C}$$

Instantiation

$$\frac{A}{A/\sigma}$$

Induction

Reflexivity Axiom

$$x = x$$

Equality Axiom

$$x_1 = y_1 \rightarrow x_2 = y_2 \rightarrow x_1 = x_2 \rightarrow y_1 = y_2$$

Referential Transparency

$$x_1 = y_1 \rightarrow \dots \rightarrow x_n = y_n \rightarrow f(x_1, \dots, x_n) = f(y_1, \dots, y_n)$$

Beta Reduction

$$((\lambda x_1 \dots x_n . \beta) t_1 \dots t_n) = \beta/[x_1 \leftarrow t_1, \dots, x_n \leftarrow t_n]$$

Base Evaluation

$$\text{e.g., } 1 + 2 = 3$$

Lisp Axioms

$$\text{e.g., } \text{consp}(\text{cons}(x, y)) = t$$

Milawa's logic

Prop. Schema $\frac{}{\neg A \vee A}$

Contraction $\frac{A \vee A}{A}$

Expansion $\frac{A}{B \vee A}$

Associativity $\frac{A \vee (B \vee C)}{(A \vee B) \vee C}$

Cut $\frac{A \vee B \quad \neg A \vee C}{B \vee C}$

Instantiation $\frac{A}{A/\sigma}$

w.r.t. ordinals up to ϵ_0

Induction

Reflexivity Axiom

$$x = x$$

Equality Axiom

$$x_1 = y_1 \rightarrow x_2 = y_2 \rightarrow x_1 = x_2 \rightarrow y_1 = y_2$$

Referential Transparency

$$x_1 = y_1 \rightarrow \dots \rightarrow x_n = y_n \rightarrow f(x_1, \dots, x_n) = f(y_1, \dots, y_n)$$

Beta Reduction

$$((\lambda x_1 \dots x_n . \beta) t_1 \dots t_n) = \beta/[x_1 \leftarrow t_1, \dots, x_n \leftarrow t_n]$$

Base Evaluation

$$\text{e.g., } 1 + 2 = 3$$

Lisp Axioms

$$\text{e.g., } \text{consp}(\text{cons}(x, y)) = t$$

Milawa's logic

Prop. Schema $\frac{}{\neg A \vee A}$

Contraction $\frac{A \vee A}{A}$

Expansion $\frac{A}{B \vee A}$

Associativity $\frac{A \vee (B \vee C)}{(A \vee B) \vee C}$

Cut $\frac{A \vee B \quad \neg A \vee C}{B \vee C}$

Instantiation $\frac{A}{A/\sigma}$

w.r.t. ordinals up to ϵ_0
 Induction

Reflexivity Axiom
 $x = x$

Equality Axiom
 $x_1 = y_1 \rightarrow x_2 = y_2 \rightarrow x_1 = x_2 \rightarrow y_1 = y_2$

Referential Transparency
 $x_1 = y_1 \rightarrow \dots \rightarrow x_n = y_n \rightarrow f(x_1, \dots, x_n) = f(y_1, \dots, y_n)$

Beta Reduction
 $((\lambda x \dots x \beta) t \dots t) \equiv \beta / [x \leftarrow t \dots x \leftarrow t]$

evaluation of any lisp primitive applied to constants

Base Evaluation
 e.g., $1+2 = 3$

Lisp Axioms
 e.g., $cons(cons(x, y)) = t$

Milawa's logic

Prop. Schema $\frac{}{\neg A \vee A}$

Contraction $\frac{A \vee A}{A}$

Expansion $\frac{A}{B \vee A}$

Associativity $\frac{A \vee (B \vee C)}{(A \vee B) \vee C}$

Cut $\frac{A \vee B \quad \neg A \vee C}{B \vee C}$

Instantiation $\frac{A}{A/\sigma}$

w.r.t. ordinals up to ϵ_0

Induction

Reflexivity Axiom
 $x = x$

Equality Axiom
 $x_1 = y_1 \rightarrow x_2 = y_2 \rightarrow x_1 = x_2 \rightarrow y_1 = y_2$

Referential Transparency
 $x_1 = y_1 \rightarrow \dots \rightarrow x_n = y_n \rightarrow f(x_1, \dots, x_n) = f(y_1, \dots, y_n)$

Beta Reduction
 $((\lambda x. x \beta) t) \equiv \beta / [x \leftarrow t]$

evaluation of any lisp primitive applied to constants

Base Evaluation
e.g., $1+2 = 3$

Lisp Axioms
e.g., $cons(cons(x, y)) = t$

56 axioms describing properties of Lisp primitives

Trusting Milawa

Milawa is trustworthy if:

- logic is sound
- core implements the logic correctly
- runtime executes the core correctly

If the above are proved, then Milawa could be “the most trustworthy theorem prover”.

Outline

Part 1: **Milawa**

Part 2: **Its new verified runtime**

Part 3: **Mini-demos, measurements**

Outline

Part 1: **Milawa**

Part 2: **Its new verified runtime** 

Part 3: **Mini-demos, measurements**

Requirements on runtime

Milawa uses a subset of Common Lisp which

is for most part **first-order pure functions** over
natural numbers, symbols and conses,

uses primitives: `cons car cdr consp natp symbolp
equal + - < symbol-< if`

macros: `or and list let let* cond
first second third fourth fifth`

and a simple form of lambda-applications.

(Lisp subset defined on later slide.)

Requirements on runtime

...but Milawa also

- uses **destructive updates**, hash tables
- prints status messages, **timing data**
- uses Common Lisp's **checkpoints**
- forces function **compilation**
- makes **dynamic function calls**
- can produce **runtime errors**

(Lisp subset defined on later slide.)

Requirements on runtime

...but Milawa also

- ~~uses destructive updates, hash tables~~
- ~~prints status messages, timing data~~
- ~~uses Common Lisp's checkpoints~~
- forces function compilation
- makes dynamic function calls
- can produce runtime errors

(Lisp subset defined on later slide.)

Requirements on runtime

...but Milawa also

- ~~uses destructive updates, hash tables~~
 - ~~prints status messages, timing data~~
 - ~~uses Common Lisp's checkpoints~~
 - forces function compilation
 - makes dynamic function calls
 - can produce runtime errors
- } not necessary
- } runtime must support

(Lisp subset defined on later slide.)

Runtime must scale

Designed to scale:

Runtime must scale

Designed to scale:

- dynamic compilation
 - ▶ functions compile to native code

Runtime must scale

Designed to scale:

- dynamic compilation
 - ▶ functions compile to native code
- target 64-bit x86 for heap capacity
 - ▶ space for 2^{31} (2 billion) cons cells (16 GB)

Runtime must scale

Designed to scale:

- dynamic compilation
 - ▶ functions compile to native code
- target 64-bit x86 for heap capacity
 - ▶ space for 2^{31} (2 billion) cons cells (16 GB)
- efficient scannerless parsing + abbreviations
 - ▶ must cope with 4 gigabyte input

Runtime must scale

Designed to scale:

- dynamic compilation
 - ▶ functions compile to native code
- target 64-bit x86 for heap capacity
 - ▶ space for 2^{31} (2 billion) cons cells (16 GB)
- efficient scannerless parsing + abbreviations
 - ▶ must cope with 4 gigabyte input
- graceful exits in all circumstances
 - ▶ allowed to run out of space, but must report it

Constructing the new runtime

Conventional method?

1. write simple code
2. then prove it correct

Method used:

1. write approximately correct algorithm implementation
2. start a verification proof
3. iteratively tweak the code and the proof

Constructing the new runtime

Step-by-step:

1. specified input language: syntax & semantics
2. verified necessary algorithms, e.g.
 - compilation from source to bytecode
 - parsing and printing of s-expressions
 - copying garbage collection
3. proved refinements from algorithms to x86 code
4. plugged together to form read-eval-print loop

AST of input language

<i>term</i>	::=	Const <i>sexp</i>	<i>sexp</i>	::=	Val <i>num</i>
		Var <i>string</i>			Sym <i>string</i>
		App <i>func</i> (<i>term</i> list)			Dot <i>sexp sexp</i>
		If <i>term term term</i>			
		LambdaApp (<i>string</i> list) <i>term</i> (<i>term</i> list)			
		Or (<i>term</i> list)			
		And (<i>term</i> list)			(macro)
		List (<i>term</i> list)			(macro)
		Let ((<i>string</i> × <i>term</i>) list) <i>term</i>			(macro)
		LetStar ((<i>string</i> × <i>term</i>) list) <i>term</i>			(macro)
		Cond ((<i>term</i> × <i>term</i>) list)			(macro)
		First <i>term</i> Second <i>term</i> Third <i>term</i>			(macro)
		Fourth <i>term</i> Fifth <i>term</i>			(macro)
<i>func</i>	::=	Define Print Error Funcall			
		PrimitiveFun <i>primitive</i> Fun <i>string</i>			
<i>primitive</i>	::=	Equal Symbolp SymbolLess			
		Consp Cons Car Cdr			
		Natp Add Sub Less			

AST of input language

Example of semantics for macros:

$$\frac{(\text{App } (\text{PrimitiveFun } \text{Car}) [x], \text{env}, k, io) \xrightarrow{\text{ev}} (\text{ans}, \text{env}', k', io')}{(\text{First } x, \text{env}, k, io) \xrightarrow{\text{ev}} (\text{ans}, \text{env}', k', io')}$$

		List (<i>term list</i>)	(macro)
		Let ((<i>string</i> × <i>term</i>) list) <i>term</i>	(macro)
		LetStar ((<i>string</i> × <i>term</i>) list) <i>term</i>	(macro)
		Cond ((<i>term</i> × <i>term</i>) list)	(macro)
		First <i>term</i> Second <i>term</i> Third <i>term</i>	(macro)
		Fourth <i>term</i> Fifth <i>term</i>	(macro)
<i>func</i>	::=	Define Print Error Funcall	
		PrimitiveFun <i>primitive</i> Fun <i>string</i>	
<i>primitive</i>	::=	Equal Symbolp SymbolLess	
		Consp Cons Car Cdr	
		Natp Add Sub Less	

compile: AST \rightarrow bytecode list

<i>bytecode</i>	::=	Pop	pop one stack element
		PopN <i>num</i>	pop <i>n</i> stack elements
		PushVal <i>num</i>	push a constant number
		PushSym <i>string</i>	push a constant symbol
		LookupConst <i>num</i>	push the <i>n</i> th constant from system state
		Load <i>num</i>	push the <i>n</i> th stack element
		Store <i>num</i>	overwrite the <i>n</i> th stack element
		DataOp <i>primitive</i>	add, subtract, car, cons, ...
		Jump <i>num</i>	jump to program point <i>n</i>
		JumpIfNil <i>num</i>	conditionally jump to <i>n</i>
		DynamicJump	jump to location given by stack top
		Call <i>num</i>	static function call (faster)
		DynamicCall	dynamic function call (slower)
		Return	return to calling function
		Fail	signal a runtime error
		Print	print an object to stdout
		Compile	compile a function definition

How do we get compilation to x86?

We have verified compilation algorithm:

compile: AST \rightarrow bytecode list

but compiler must produce real x86 code....

How do we get compilation to x86?

We have verified compilation algorithm:

compile: AST \rightarrow bytecode list

but compiler must produce real x86 code....

Solution:

- bytecode is represented by numbers in memory that are x86 machine code
- we prove that jumping to the memory location of the bytecode executes it

How do we get compilation to x86?

Treating code as data:

$$\forall p \ c \ q. \ \{p\} \ c \ \{q\} = \{p * \text{code } c\} \ \emptyset \ \{q * \text{code } c\}$$

(POPL'10)

Solution:

- bytecode is represented by numbers in memory that are x86 machine code
- we prove that jumping to the memory location of the bytecode executes it

I/O and efficient parsing

Jitawa implements a read-eval-print loop:

Use of external **C routines** adds assumptions to proof:

- reading next string from stdin
- printing null-terminated string to stdout

I/O and efficient parsing

Jitawa implements a read-eval-print loop:

Use of external **C routines** adds assumptions to proof:

- reading next string from stdin
- printing null-terminated string to stdout

An efficient **s-expression parser** (and **printer**) is proved, which deals with abbreviations:

```
(append (cons (cons a b) c)
        (cons (cons a b) c))
```

```
(append #1=(cons (cons a b) c)
        #1#)
```

Read-eval-print loop

- Result of reading **lazily**, writing **eagerly**
- Eval = **compile then jump-to-compiled-code**
- Specification: read-eval-print until end of input

Correctness theorem

Top-level correctness theorem:

$$\{ \text{init_state } io * \text{pc } p * \langle \text{terminates_for } io \rangle \}$$
$$p : \text{code_for_entire_jitawa_implementation}$$
$$\{ \text{error_message } \vee \exists io'. \langle ([], io) \xrightarrow{\text{exec}} io' \rangle * \text{final_state } io' \}$$

Correctness theorem

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

ness theorem:

$$\{ \text{init_state } io * \text{pc } p * \langle \text{terminates_for } io \rangle \}$$
$$p : \text{code_for_entire_jitawa_implementation}$$
$$\{ \text{error_message } \vee \exists io'. \langle ([], io) \xrightarrow{\text{exec}} io' \rangle * \text{final_state } io' \}$$

Correctness theorem

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

ness theorem:

$$\{ \text{init_state } io * \text{pc } p * \langle \text{terminates_for } io \rangle \}$$
$$p : \text{code_for_entire_jitawa_implementation}$$
$$\{ \text{error_message } \vee \exists io'. \langle ([], io) \xrightarrow{\text{exec}} io' \rangle * \text{final_state } io' \}$$

Each execution is allowed to fail with an error message.

Correctness theorem

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

Correctness theorem:

$$\{ \text{init_state } io * \text{pc } p * \langle \text{terminates_for } io \rangle \}$$

p : code_for_entire_jitawa_implementation

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

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.

Correctness theorem

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

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

$\{ \text{init_state } io * \text{pc } p * \langle \text{terminates_for } io \rangle \}$

$p : \text{code_for_entire_jitawa_implementation}$

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

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.

Correctness theorem

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

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

$\{ \text{init_state } io * \text{pc } p * \langle \text{terminates_for } io \rangle \}$

$p : \text{code_for_entire_jitawa_implementation}$ list of numbers

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

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.

Verified code

```
$ cat verified_code.s
```

```
/* Machine code automatically extracted from a HOL4 theorem. */
```

```
/* The code consists of 7423 instructions (31840 bytes). */
```

```
.byte 0x48, 0x8B, 0x5F, 0x18
```

```
.byte 0x4C, 0x8B, 0x7F, 0x10
```

```
.byte 0x48, 0x8B, 0x47, 0x20
```

```
.byte 0x48, 0x8B, 0x4F, 0x28
```

```
.byte 0x48, 0x8B, 0x57, 0x08
```

```
.byte 0x48, 0x8B, 0x37
```

```
.byte 0x4C, 0x8B, 0x47, 0x60
```

```
.byte 0x4C, 0x8B, 0x4F, 0x68
```

```
.byte 0x4C, 0x8B, 0x57, 0x58
```

```
.byte 0x48, 0x01, 0xC1
```

```
.byte 0xC7, 0x00, 0x04, 0x4E, 0x49, 0x4C
```

```
.byte 0x48, 0x83, 0xC0, 0x04
```

```
.byte 0xC7, 0x00, 0x02, 0x54, 0x06, 0x51
```

```
.byte 0x48, 0x83, 0xC0, 0x04
```

```
...
```

Verified code

```
$ cat verified_code.s
```

```
/* Machine code automatically extracted from a HOL4 theorem. */  
/* The code consists of 7423 instructions (31840 bytes). */
```

```
.byte 0x48, 0x8B, 0x5F, 0x18  
.byte 0x4C, 0x8B, 0x7F, 0x10  
.byte 0x48, 0x8B, 0x47, 0x20  
.byte 0x48, 0x8B, 0x4F, 0x28  
.byte 0x48, 0x8B, 0x57, 0x08  
.byte 0x48, 0x8B, 0x37
```

How is this verified binary produced?

Demo: proof-producing synthesis (TPHOLs'09)

```
.byte 0xC7, 0x00, 0x02, 0x54, 0x06, 0x51  
.byte 0x48, 0x83, 0xC0, 0x04  
...
```


Outline

Part 1: **Milawa**

Part 2: **Its new verified runtime**

Part 3: **Mini-demos, measurements**

Outline

Part 1: **Milawa**

Part 2: **Its new verified runtime**

Part 3: **Mini-demos, measurements**



Running Milawa on Jitawa

Running Milawa's 4-gigabyte bootstrap process:

CCL	16 hours	
SBCCL	22 hours	
Jitawa	128 hours	(8x slower than CCL)

Running Milawa on Jitawa

Running Milawa's 4-gigabyte bootstrap process:

CCL	16 hours	
SBCCL	22 hours	
Jitawa	128 hours	(8x slower than CCL)

Jitawa's compiler performs almost no optimisations.

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

Quirky behaviour

DEMO

Jitawa mimics an interpreter's behaviour

- to **hide** the fact that **compilation** occurs
- to keep semantics as **simple** as possible
- to facilitate **future work** (e.g. verify Milawa's core)

Quirky behaviour

DEMO

Jitawa mimics an interpreter's behaviour

- to **hide** the fact that **compilation** occurs
- to keep semantics as **simple** as possible
- to facilitate **future work** (e.g. verify Milawa's core)

Consequences:

- compiler must turn **undefined functions, bad arity and unknown variables** into runtime checks/fails.
- **mutual recursion** is free!

Conclusions

Summary

- new verified runtime
- implements clean Lisp language
- scales and is able to host Milawa theorem prover

Next year?

- Milawa proved sound down to the machine code which runs it?

Conclusions

Summary

- new verified runtime
- implements clean Lisp language
- scales and is able to host Milawa theorem prover

Next year?

- Milawa proved sound down to the machine code which runs it?

Questions?