

The ACL2 Theorem Prover and How It Came to be Used in Industry

(A Sprint through 45 Years of Verification History)

J Strother Moore
Department of Computer Science
University of Texas at Austin


Boyer-Moore-Kaufmann Project

Edinburgh Pure Lisp Theorem Prover [*BM, 1973*]

A Computational Logic [*BM, 1979*]

NQTHM [*BM, 1981*]

ACL2 [*BM, KM 1989–present*]



Vision: To build a practical verification system: axiomatize a functional programming language, build a theorem prover for it, and use them to model and verify computational artifacts.

ACL2

A **C**omputational **L**ogic for
Applicative **C**ommon **L**isp

A fully integrated verification environment for a practical applicative subset of an ANSI standard programming language

{kaufmann,moore}@cs.utexas.edu

<http://www.cs.utexas.edu/users/moore/acl2>

I can't explain ACL2 in the time we have, I can only highlight some capabilities.

History of Theorems Proved (by Our Provers)

simple list processing

academic math and CS

breakthrough commercial applications

routine industrial use



A Few Axioms

- $t \neq \text{nil}$
- $x = \text{nil} \rightarrow (\text{if } x \ y \ z) = z$
- $x \neq \text{nil} \rightarrow (\text{if } x \ y \ z) = y$
- $(\text{car } (\text{cons } x \ y)) = x$
- $(\text{cdr } (\text{cons } x \ y)) = y$
- $(\text{endp } \text{nil}) = t$
- $(\text{endp } (\text{cons } x \ y)) = \text{nil}$

A Common Abbreviation

`'(1 2 3)`

`=`

`(cons 1 '(2 3))`

`=`

`...`

`(cons 1 (cons 2 (cons 3 nil)))`

Definition – List Concatenation

```
(defun ap (x y)
  (if (endp x)
      y
      (cons (car x)
            (ap (cdr x) y))))
```

Definition – List Concatenation

```
(defun ap (x y)
  (if (endp x)
      y
      (cons (car x)
            (ap (cdr x) y))))
```

Thm: $(\text{ap } \text{nil } y) = y$

Thm: $(\text{ap } (\text{cons } u \ v) \ y) = (\text{cons } u \ (\text{ap } v \ y))$

Definition – List Concatenation

```
(defun ap (x y)
  (if (endp x)
      y
      (cons (car x)
            (ap (cdr x) y))))
```

Thm: $(\text{ap } \text{nil } y) = y$

Thm: $(\text{ap } (\text{cons } u \ v) \ y) = (\text{cons } u \ (\text{ap } v \ y))$

$(\text{ap } '(1 \ 2 \ 3) \ '(4 \ 5 \ 6)) =$
 $(\text{ap } '(1 \ 2 \ 3) \ '(4 \ 5 \ 6))$

Definition – List Concatenation

```
(defun ap (x y)
  (if (endp x)
      y
      (cons (car x)
            (ap (cdr x) y))))
```

Thm: $(\text{ap } \text{nil } y) = y$

Thm: $(\text{ap } (\text{cons } u \ v) \ y) = (\text{cons } u \ (\text{ap } v \ y))$

```
(ap '(1 2 3) '(4 5 6)) =
(ap (cons 1 (cons 2 (cons 3 nil))) '(4 5 6))
```

Definition – List Concatenation

```
(defun ap (x y)
  (if (endp x)
      y
      (cons (car x)
            (ap (cdr x) y))))
```

Thm: $(\text{ap } \text{nil } y) = y$

Thm: $(\text{ap } (\text{cons } u \ v) \ y) = (\text{cons } u \ (\text{ap } v \ y))$

```
(ap '(1 2 3) '(4 5 6)) =
(cons 1 (ap (cons 2 (cons 3 nil)) '(4 5 6)))
```

Definition – List Concatenation

```
(defun ap (x y)
  (if (endp x)
      y
      (cons (car x)
            (ap (cdr x) y))))
```

Thm: $(\text{ap } \text{nil } y) = y$

Thm: $(\text{ap } (\text{cons } u \ v) \ y) = (\text{cons } u \ (\text{ap } v \ y))$

```
(ap '(1 2 3) '(4 5 6)) =
(cons 1 (cons 2 (ap (cons 3 nil) '(4 5 6))))
```

Definition – List Concatenation

```
(defun ap (x y)
  (if (endp x)
      y
      (cons (car x)
            (ap (cdr x) y))))
```

Thm: $(\text{ap } \text{nil } y) = y$

Thm: $(\text{ap } (\text{cons } u \ v) \ y) = (\text{cons } u \ (\text{ap } v \ y))$

```
(ap '(1 2 3) '(4 5 6)) =
(cons 1 (cons 2 (cons 3 (ap nil '(4 5 6)))))
```

Definition – List Concatenation

```
(defun ap (x y)
  (if (endp x)
      y
      (cons (car x)
            (ap (cdr x) y))))
```

Thm: $(\text{ap } \text{nil } y) = y$

Thm: $(\text{ap } (\text{cons } u \ v) \ y) = (\text{cons } u \ (\text{ap } v \ y))$

```
(ap '(1 2 3) '(4 5 6)) =
(cons 1 (cons 2 (cons 3 '(4 5 6))))
```

Definition – List Concatenation

```
(defun ap (x y)
  (if (endp x)
      y
      (cons (car x)
            (ap (cdr x) y))))
```

Thm: $(\text{ap } \text{nil } y) = y$

Thm: $(\text{ap } (\text{cons } u \ v) \ y) = (\text{cons } u \ (\text{ap } v \ y))$

```
(ap '(1 2 3) '(4 5 6)) =
(cons 1 (cons 2 '(3 4 5 6)))
```

Definition – List Concatenation

```
(defun ap (x y)
  (if (endp x)
      y
      (cons (car x)
            (ap (cdr x) y)))))
```

Thm: $(\text{ap } \text{nil } y) = y$

Thm: $(\text{ap } (\text{cons } u \ v) \ y) = (\text{cons } u \ (\text{ap } v \ y))$

```
(ap '(1 2 3) '(4 5 6)) =
(cons 1 '(2 3 4 5 6))
```


Definition – List Concatenation

```
(defun ap (x y)
  (if (endp x)
      y
      (cons (car x)
            (ap (cdr x) y))))
```

Thm: $(\text{ap } \text{nil } y) = y$

Thm: $(\text{ap } (\text{cons } u \ v) \ y) = (\text{cons } u \ (\text{ap } v \ y))$

```
(ap '(1 2 3) '(4 5 6)) =
'(1 2 3 4 5 6)
```

Theorem ap is associative (1971)

$$\forall a \forall b \forall c : \text{ap}(\text{ap}(a,b),c) = \text{ap}(a,\text{ap}(b,c)).$$

(equal (ap (ap a b) c)
 (ap a (ap b c)))

```
(equal (ap (ap a b) c)
       (ap a (ap b c)))
```

(equal (ap (ap a b) c)
 (ap a (ap b c))))

Proof: by induction on a.

(equal (ap (ap a b) c)
 (ap a (ap b c))))

Proof: by induction on a.

Base Case: (endp a).

(equal (ap (ap a b) c)
 (ap a (ap b c))))

(equal (ap (ap a b) c)
 (ap a (ap b c))))

Proof: by induction on a.

Base Case: (endp a).

(equal (ap b c)
 (ap a (ap b c))))

(equal (ap (ap a b) c)
 (ap a (ap b c))))

Proof: by induction on a.

Base Case: (endp a).

(equal (ap b c)
 (ap a (ap b c)))

(equal (ap (ap a b) c)
 (ap a (ap b c))))

Proof: by induction on a.

Base Case: (endp a).

(equal (ap b c)
 (ap b c))

(equal (ap (ap a b) c)
 (ap a (ap b c)))

Proof: by induction on a.

Base Case: (endp a).

(equal (ap b c)
 (ap b c))

(equal (ap (ap a b) c)
 (ap a (ap b c))))

Proof: by induction on a.

Base Case: (endp a).

T

```
(equal (ap (ap a b) c)
       (ap a (ap b c)))
```

Proof: by induction on a.

Induction Step: (not (endp a)).

```
(equal (ap (ap a b) c)
       (ap a (ap b c)))
```

(equal (ap (ap (cdr a) b) c) ; *Ind Hyp*
 (ap (cdr a) (ap b c)))

Proof: by induction on a.

Induction Step: (not (endp a)).

(equal (ap (ap a b) c)
 (ap a (ap b c)))

(equal (ap (ap (cdr a) b) c) ; *Ind Hyp*
(ap (cdr a) (ap b c)))

Proof: by induction on a.

Induction Step: (not (endp a)).

(equal (ap (cons (car a)
(ap (cdr a) b)) c)
(ap a (ap b c)))

(equal (ap (ap (cdr a) b) c) ; *Ind Hyp*
(ap (cdr a) (ap b c)))

Proof: by induction on a.

Induction Step: (not (endp a)).

(equal (ap (cons (car a)
(ap (cdr a) b)) c)
(ap a (ap b c)))

(equal (ap (ap (cdr a) b) c) ; *Ind Hyp*
(ap (cdr a) (ap b c)))

Proof: by induction on a.

Induction Step: (not (endp a)).

(equal (cons (car a)
(ap (ap (cdr a) b) c))
(ap a (ap b c)))

(equal (ap (ap (cdr a) b) c) ; *Ind Hyp*
 (ap (cdr a) (ap b c)))

Proof: by induction on a.

Induction Step: (not (endp a)).

(equal (cons (car a)
 (ap (ap (cdr a) b) c))
 (ap a (ap b c)))

(equal (ap (ap (cdr a) b) c) ; *Ind Hyp*
 (ap (cdr a) (ap b c)))

Proof: by induction on a.

Induction Step: (not (endp a)).
(equal (cons (car a)
 (ap (ap (cdr a) b) c))
 (cons (car a)
 (ap (cdr a) (ap b c)))))

(equal (ap (ap (cdr a) b) c) ; *Ind Hyp*
(ap (cdr a) (ap b c)))

Proof: by induction on a.

Induction Step: (not (endp a)).
(equal (cons (car a)
 (ap (ap (cdr a) b) c))
(cons (car a)
 (ap (cdr a) (ap b c))))

(equal (ap (ap (cdr a) b) c) ; *Ind Hyp*
 (ap (cdr a) (ap b c)))

Proof: by induction on a.

Induction Step: (not (endp a)).

(equal
 (ap (ap (cdr a) b) c)
 (ap (cdr a) (ap b c)))

(equal (ap (ap (cdr a) b) c) ; *Ind Hyp*
 (ap (cdr a) (ap b c)))

Proof: by induction on a.

Induction Step: (not (endp a)).
(equal (ap (ap (cdr a) b) c)
 (ap (cdr a) (ap b c)))

(equal (ap (ap (cdr a) b) c) ; *Ind Hyp*
 (ap (cdr a) (ap b c)))

Proof: by induction on a.

Induction Step: (not (endp a)).
(equal (ap (ap (cdr a) b) c)
 (ap (cdr a) (ap b c)))

(equal (ap (ap a b) c)
 (ap a (ap b c))))

Proof: by induction on a.

Induction Step: (not (endp a)).

T

(equal (ap (ap a b) c)
 (ap a (ap b c))))

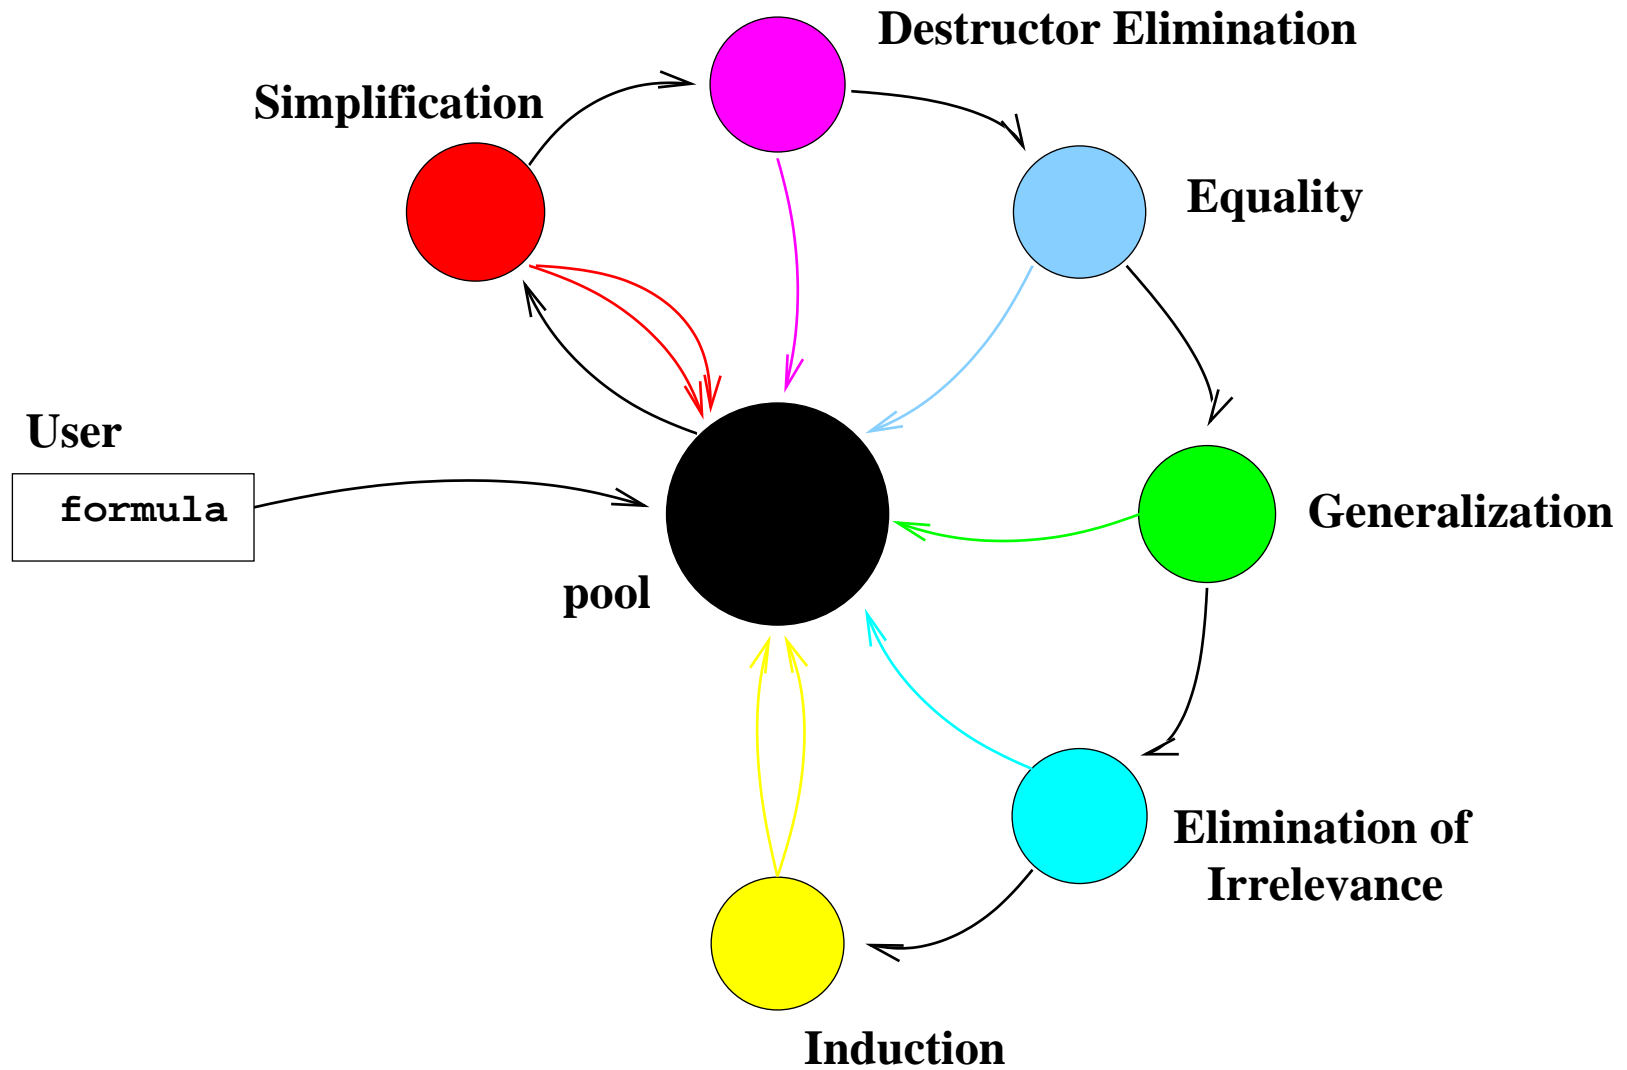
Proof: by induction on a.

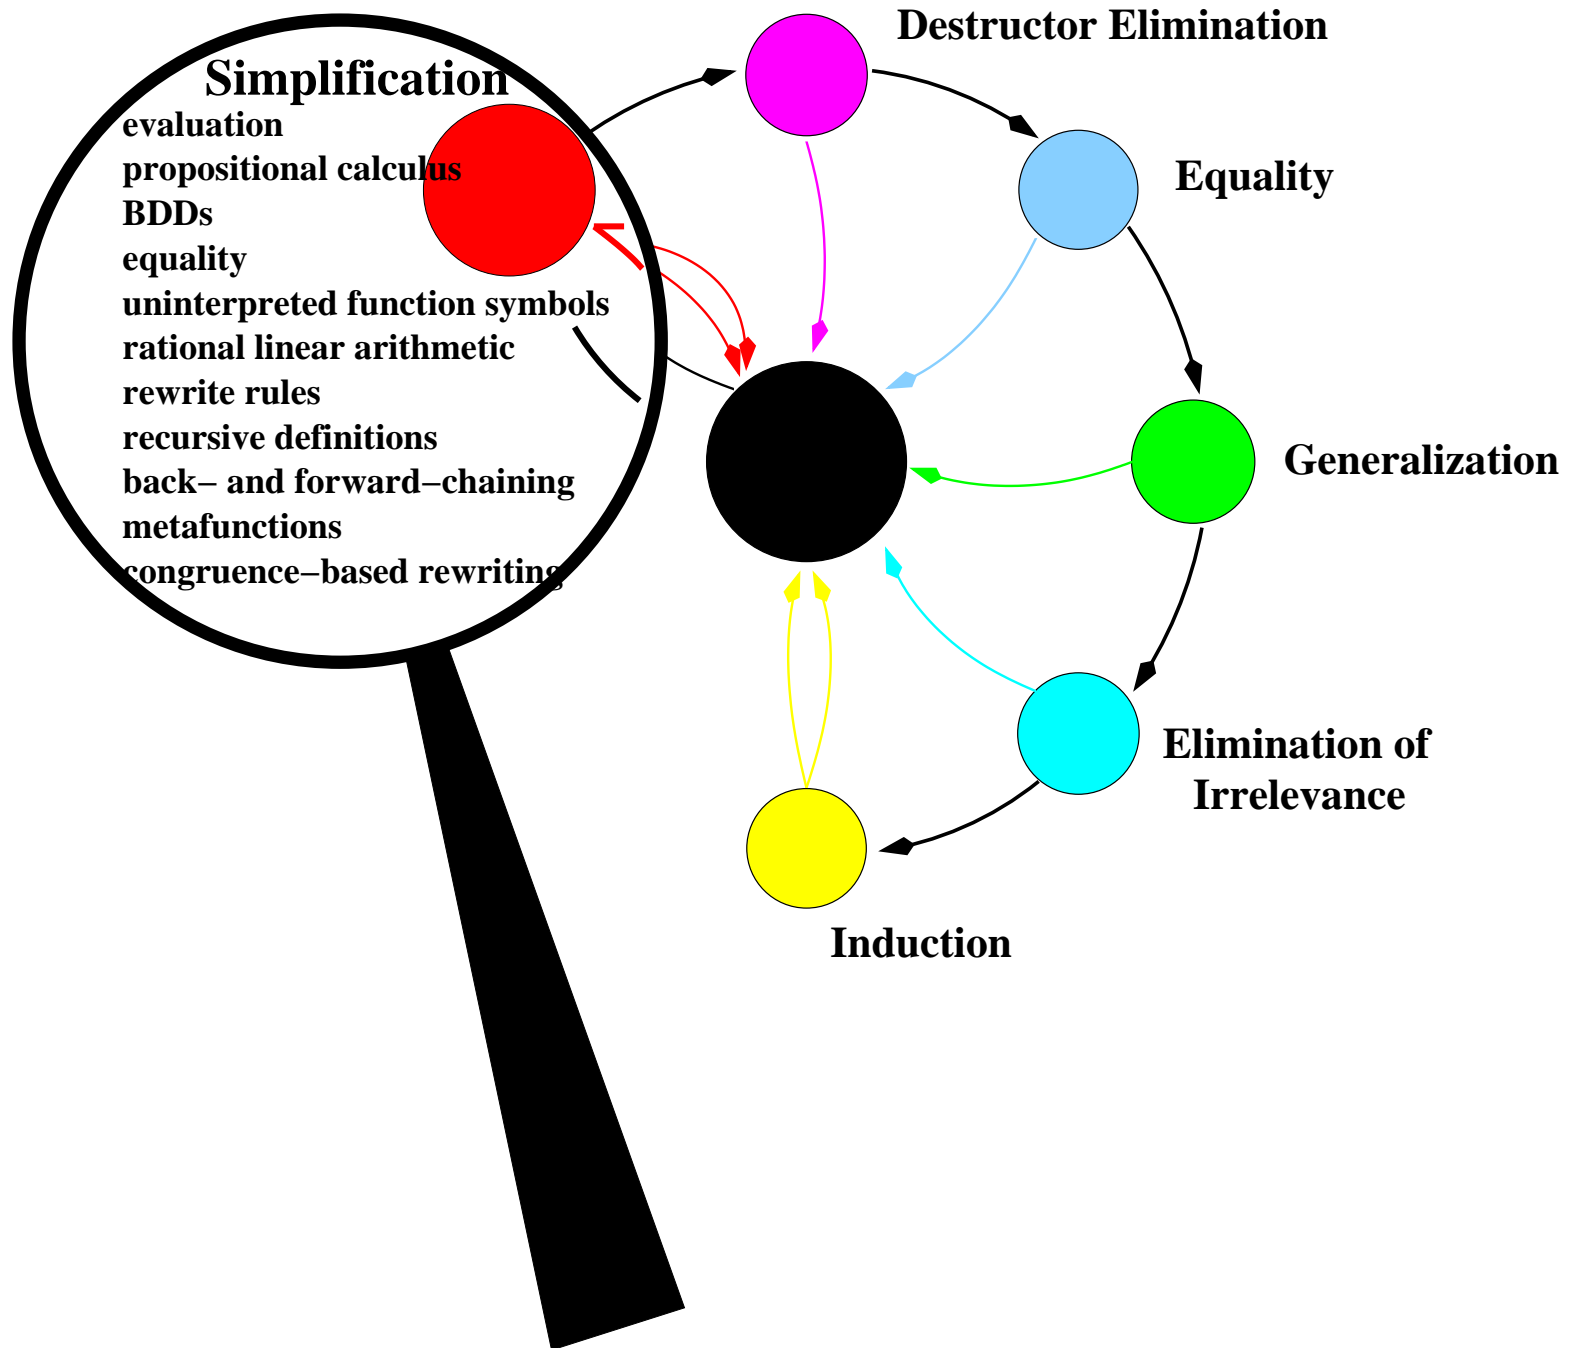
Q.E.D.


ACL2 Demo 1


Notes: The four demos in this talk are with ACL2 even though many of the proofs in the first few demos were first done with earlier Boyer-Moore provers.


The ACL2 image used here has been configured to support these four demos.















ACL2 Community Books (2017)



<https://github.com/acl2/acl2/books/>
contains $\sim 6,000$ user-supplied books.

ACL2 Demo 2

History of Theorems Proved

simple list processing

academic math and CS

breakthrough commercial
applications

routine industrial use



Academic Math (Nqthm, 1980s)

- undecidability of the halting problem
(18 lemmas)
- invertibility of RSA encryption
(172 lemmas)
- Gauss' law of quadratic reciprocity [Russinoff]
(348 lemmas)
- Gödel's First Incompleteness Theorem [Shankar]
(1741 lemmas)

Academic CS (Nqthm, 1980s)

- The CLI Verified Stack:
 - microprocessor: gates to machine code [Hunt]
 - assembler-linker-loader (3326 lemmas) [Moore]
 - compilers [Young, Flatau]
 - operating system [Bevier]
 - applications [Wilding]

These theorems guarantee that a property proved about an app holds when it is compiled, assembled, linked, loaded, and run on the gate-level machine.

ACL2 Demo 3

Seeing Nqthm struggle with “large” models like Piton and other components of the CLI verified stack convinced us to re-implement it with scaling in mind.

Thus was born **ACL2** (1989).

History of Theorems Proved

simple list processing

academic math and CS

breakthrough commercial
applications

routine industrial use



An elusive circuitry error is causing a chip used in millions of computers to generate inaccurate results

— *NY Times*, “Circuit Flaw Causes Pentium Chip to Miscalculate, Intel Admits,” Nov 11, 1994

Intel Corp. last week took a \$475 million write-off to cover costs associated with the divide bug in the Pentium microprocessor's floating-point unit — *EE Times, Jan 23, 1995*

IEEE 754 Floating Point Standard

Elementary operations are to be performed as though the infinitely precise (standard mathematical) operation were performed and then the result rounded to the indicated precision.

AMD K5 Algorithm FDIV($p, d, mode$)

1.	$sd_0 = \text{lookup}(d)$	[exact	17	8]
2.	$d_r = d$	[away	17	32]
3.	$sdd_0 = sd_0 \times d_r$	[away	17	32]
4.	$sd_1 = sd_0 \times \text{comp}(sdd_0, 32)$	[trunc	17	32]
5.	$sdd_1 = sd_1 \times d_r$	[away	17	32]
6.	$sd_2 = sd_1 \times \text{comp}(sdd_1, 32)$	[trunc	17	32]
...	... =		
29.	$q_3 = sd_2 \times ph_3$	[trunc	17	24]
30.	$qq_2 = q_2 + q_3$	[sticky	17	64]
31.	$qq_1 = qq_2 + q_1$	[sticky	17	64]
32.	$fdiv = qq_1 + q_0$	<i>mode</i>		

Using the Reciprocal

$$\begin{array}{r}
 36. \\
 + \quad -17 \\
 + \quad .0034 \\
 + \quad \underline{-000066} \\
 \hline
 35.833334 \\
 12 \overline{) 430.000000} \\
 \underline{432.} \\
 -2. \\
 \underline{-2.04} \\
 .04 \\
 \underline{.0408} \\
 - .0008 \\
 \underline{- .000792} \\
 - .000008
 \end{array}$$

Reciprocal Calculation:

$$1/12 = 0.083\overline{3} \approx 0.083 = sd_2$$

Quotient Digit Calculation:

$$0.083 \times 430.0000 = 35.690000 \approx 36.000000 = q_0$$

$$0.083 \times -2.0000 = -.166000 \approx -.170000 = q_1$$


$$0.083 \times .0400 = .003320 \approx .003400 = q_2$$

$$0.083 \times -.0008 = -.000664 \approx -.00067 = q_3$$

Summation of Quotient Digits:

$$q_0 + q_1 + q_2 + q_3 = 35.833333$$

Computing the Reciprocal



top 8 bits of d	approx inverse	top 8 bits of d	approx inverse	top 8 bits of d	approx inverse	top 8 bits of d	approx inverse
1.000000 ₂	0.1111111 ₂	1.010000 ₂	0.1100110 ₂	1.100000 ₂	0.1010101 ₂	1.110000 ₂	0.1001001 ₂
1.000001 ₂	0.1111101 ₂	1.010001 ₂	0.1100101 ₂	1.100001 ₂	0.1010100 ₂	1.110001 ₂	0.1001000 ₂
1.000010 ₂	0.1111101 ₂	1.010010 ₂	0.1100101 ₂	1.100010 ₂	0.1010100 ₂	1.110010 ₂	0.1001000 ₂
1.000011 ₂	0.1111100 ₂	1.010011 ₂	0.1100100 ₂	1.100011 ₂	0.1010100 ₂	1.110011 ₂	0.1001000 ₂
1.000100 ₂	0.1111011 ₂	1.010010 ₂	0.1100011 ₂	1.100100 ₂	0.1010011 ₂	1.110010 ₂	0.1000111 ₂
1.000101 ₂	0.1111010 ₂	1.010010 ₂	0.1100011 ₂	1.100101 ₂	0.1010011 ₂	1.110010 ₂	0.1000111 ₂
1.000110 ₂	0.1111010 ₂	1.010011 ₂	0.1100010 ₂	1.100110 ₂	0.1010010 ₂	1.110011 ₂	0.1000110 ₂
1.000111 ₂	0.1111001 ₂	1.010011 ₂	0.1100010 ₂	1.100111 ₂	0.1010010 ₂	1.110011 ₂	0.1000110 ₂
1.000100 ₂	0.1111000 ₂	1.010100 ₂	0.1100001 ₂	1.100100 ₂	0.1010001 ₂	1.110100 ₂	0.1000101 ₂
1.000101 ₂	0.1110111 ₂	1.010100 ₂	0.1100001 ₂	1.100101 ₂	0.1010001 ₂	1.110100 ₂	0.1000100 ₂
1.000101 ₂	0.1110110 ₂	1.010101 ₂	0.1100000 ₂	1.100101 ₂	0.1010001 ₂	1.110101 ₂	0.1000100 ₂
...
1.001011 ₂	0.1101101 ₂	1.011011 ₂	0.1011010 ₂	1.101011 ₂	0.1001100 ₂	1.111011 ₂	0.1000010 ₂
1.001011 ₂	0.1101100 ₂	1.011011 ₂	0.1011001 ₂	1.101011 ₂	0.1001100 ₂	1.111011 ₂	0.1000010 ₂
1.001100 ₂	0.1101011 ₂	1.011100 ₂	0.1011001 ₂	1.101100 ₂	0.1001011 ₂	1.111100 ₂	0.1000010 ₂
1.001100 ₂	0.1101011 ₂	1.011100 ₂	0.1011001 ₂	1.101100 ₂	0.1001011 ₂	1.111100 ₂	0.1000001 ₂
1.001101 ₂	0.1101010 ₂	1.011101 ₂	0.1011000 ₂	1.101101 ₂	0.1001010 ₂	1.111101 ₂	0.1000001 ₂
1.001101 ₂	0.1101010 ₂	1.011101 ₂	0.1011000 ₂	1.101101 ₂	0.1001010 ₂	1.111101 ₂	0.1000001 ₂
1.001101 ₂	0.1101001 ₂	1.011101 ₂	0.1010111 ₂	1.101101 ₂	0.1001010 ₂	1.111101 ₂	0.1000001 ₂
1.001110 ₂	0.1101000 ₂	1.011110 ₂	0.1010110 ₂	1.101110 ₂	0.1001010 ₂	1.111110 ₂	0.1000001 ₂
1.001110 ₂	0.1101000 ₂	1.011110 ₂	0.1010110 ₂	1.101110 ₂	0.1001010 ₂	1.111110 ₂	0.1000001 ₂
1.001110 ₂	0.1100111 ₂	1.011111 ₂	0.1010110 ₂	1.101111 ₂	0.1001001 ₂	1.111111 ₂	0.1000001 ₂
1.001111 ₂	0.1100111 ₂	1.011111 ₂	0.1010101 ₂	1.101111 ₂	0.1001001 ₂	1.111111 ₂	0.1000000 ₂
1.001111 ₂	0.1100110 ₂	1.011111 ₂	0.1010101 ₂	1.101111 ₂	0.1001001 ₂	1.111111 ₂	0.1000000 ₂

The Futility of Testing

If AMD builds this, will it work?

A bug in this design could cost AMD hundreds of millions of dollars.

On Sunway TaihuLight (93 petaflops = 93×2^{50} operations per second), testing all possible cases would take

2,726,112,523,746,722,547,161,199

$\sim 2.7 \times 10^{24}$ years.

The library of floating point lemmas and the main theorem were proved with ACL2 under the direction of two ACL2 users and the designer of the FDIV algorithm.

The proofs took 9 weeks starting from Peano's axioms.

The proofs were completed before the K5 was fabricated.

9 weeks < 2,726,112,523,746,722,547,161,199 years

The library was used in subsequent proofs.

By 1997, AMD had

- built software to translate their in-house hardware design language to ACL2
- used the tool to generate ACL2 functions modeling all the elementary floating point arithmetic on the soon-to-be fabricated AMD Athlon microprocessor

- tested the ACL2 functions by running them on AMD's standard floating-point test suite (> 100 million arithmetic problems) and compared the answers to AMD's design simulator
- proved the ACL2 functions compliant with the IEEE Standard
- found (and fixed) 3 design errors not exposed by the 100 million tests

Other Early Industrial Users of ACL2

- Motorola: DSP and microcode proofs
- AMD: floating-point on Opteron
- Rockwell-Collins: silicon JVM chip, AAMP7 crypto-box, Greenhills OS
- IBM: Power 4 FDIV and SQRT
- Sun Microsystems (via contract): Sun JVM class loader and byte-code verifier

ACL2 Demo 4

```
class Fact {
  public static int fact(int n){
    if (n>0)
      {return n*fact(n-1);}
    else return 1;
  }

  public static void main(String[] args){
    int n = Integer.parseInt(args[0], 10);
    System.out.println(fact(n));
    return;
  }
}
```


History of Theorems Proved

simple list processing

academic math and CS

breakthrough commercial
applications

routine industrial use



In 2007, Centaur Technology, Inc., challenged the ACL2 community to verify its floating-point adder:

- handles single (32-bit), double (64-bit) and extended (80-bit) additions
- pipelined to deliver 4-results per cycle
- 33,700 lines in 680 Verilog modules
- 1074 input signals (including 26 clocks) and 374 output signals

Done! After exposing and fixing one very rare bug.

(The bug occurred on exactly one pair of 80-bit inputs, i.e., 1 case of 2^{160} cases.)

ACL2 at Centaur Today

ACL2 is an indispensable part of the Centaur design process

Centaur FV team consists of 3 full-time employees and a couple of interns

Centaur has an ACL2 specification of the x86

Validated by routinely running millions of tests comparing ACL2 x86 to Intel, AMD, and Centaur hardware

The ACL2 tool-chain translates the entire Centaur design (700,000+ lines of Verilog) into a formal object in a few minutes

The translated model is validated by running millions of tests against Cadence NC Verilog and Synopsys VCS Verilog simulators

All functional-correctness proofs are re-run nightly
(on a cluster of 154 CPUs with a total of 2TB
RAM)

*“Bugs introduced today are found tonight and fixed
tomorrow.”*

Highlighted Strengths of ACL2

- adequate logical expressivity (to capture design and specs)
- adequate capacity (to manipulate multi-MB formulas)
- efficient execution (to do 100s of millions of sim runs)
- automatic proof discovery after typical design changes (to enable nightly runs)

Advantages for Centaur

- high confidence in design correctness (enabling “riskier” design changes)
- high confidence in inter-generational compatibility
- higher re-use of specifications and modules
- reduced reliance on testing
- reduced time-to-market

x86 ISA in ACL2 (Hunt and Goel)

Supports user- and system-level specifications for the x86 and may serve both to verifying (user- or system-) binary machine code and as a “build-to” spec for designers.

Performance:

user level: \sim 3.3 million ips

system level: \sim 912,000 ips

Other Ongoing Industrial Projects

- AMD (transaction protocols)
- Intel (elliptic curve crypto)
- Kestrel Institute (Android apps)
- Oracle (floating point)
- Rockwell-Collins (LLVN)
- ARM (floating point)

Industrial Wish List

- more automation (esp in lemma/defn discovery)
- faster execution speed of models in the logic
- better ways to view large formulas
- scripting capabilities
- ability to build GUIs

Things Our Industrial Users Haven't Asked For

- quantifiers
- higher-order functions
- strong typing

Our Hypothesis

The “high cost” of formal methods

– to the extent the cost is high –

is a *historical anomaly* due to the fact that virtually every project for the past 20 years has had to formalize (parts of) centuries of mathematics and decades of chip design “shop-lore”

The use of mechanized formal methods

- *decreases* time-to-market,
- *increases* reliability.

Conclusion

Mechanical reasoning systems have changed the way complex digital artifacts are built.

Complexity not an argument *against* formal methods.

It is an argument *for* formal methods.

References

Computer-Aided Reasoning: An Approach,
Kaufmann, Manolios, Moore, Kluwer Academic
Publishers, 2000.

Computer-Aided Reasoning: ACL2 Case Studies,
Kaufmann, Manolios, Moore (eds.), Kluwer
Academic Publishers, 2000.

<http://www.cs.utexas.edu/users/moore/acl2>

Extra Slides

The Formal Model of the Code

```
(defun FDIV (p d mode)
  (let*
    ((sd0 (eround (lookup d)                '(exact 17 8)))
     (dr  (eround d                          '(away 17 32)))
     (sdd0 (eround (* sd0 dr)                '(away 17 32)))
     (sd1 (eround (* sd0 (comp sdd0 32))     '(trunc 17 32)))
     (sdd1 (eround (* sd1 dr)                '(away 17 32)))
     (sd2 (eround (* sd1 (comp sdd1 32))     '(trunc 17 32)))
     ...
     (qq2 (eround (+ q2 q3)                  '(sticky 17 64)))
     (qq1 (eround (+ qq2 q1)                 '(sticky 17 64)))
     (fdiv (round (+ qq1 q0)                  mode)))
    (or (first-error sd0 dr sdd0 sd1 sdd1 ... fdiv)
        fdiv)))
```

The K5 FDIV Theorem (1200 lemmas)

“If p and d are 64-, 15⁺ floating point numbers, $d \neq 0$, and $mode$ is an IEEE rounding mode, then $FDIV(p, d, mode) = \text{round}(p/d, mode)$.”

```
(defthm FDIV-divides
  (implies (and (floating-point-numberp p 15 64)
                (floating-point-numberp d 15 64)
                (not (equal d 0))
                (rounding-modep mode))
            (equal (FDIV p d mode)
                   (round (/ p d) mode))))
```

[Moore, Kaufmann, Lynch – 1995]

A Few of the 1200 AMD Lemmas

Trunc Trunc: If $i \leq j$, then
 $\text{trunc}(\text{trunc}(x, i), j) = \text{trunc}(x, i)$.

Sticky Enough: If $mode$ is an IEEE rounding mode with size $n < i$, then
 $\text{round}(\text{sticky}(x, i + 2), mode) = \text{round}(x, mode)$.

Sticky Plus: Let $x \neq 0$, such that $\text{trunc}(x, n) = x$ and $1 + e(y) < e(x)$, and $n + e(y) - e(x) < k$.
Then $\text{sticky}(x + y, n) = \text{sticky}(x + \text{sticky}(y, k), n)$.

(Some standard hypotheses have been omitted for brevity.)

ACL2 at Centaur (cont')

Centaur uses ACL2 to build *verified* custom tools for its Verilog designers

Such tools can be used by ACL2 because of its *metafunction (reflection)* capability

Mechanized formal reasoning and theorem proving are taken for granted


Centaur's Verilog tool-chain is distributed with ACL2 and is used by Intel and Oracle

Reports on Oracle Usage

Verifying Oracle's SPARC processors with ACL2

Greg Grohoski
VP, Hardware Development
Oracle Microelectronics

May 23, 2017



Copyright © 2017, Oracle and/or its affiliates. All rights reserved.

Summer 2014

- Wrote 1st draft of ACL2 spec for
 - IEEE 754 Standard on Floating-Point Arithmetic
 - Integer division variants
- Verified
 - Floating-point implementations wrt significand
 - Integer division implementations
- Found improvements for SPARC core
 - Reduced look-up tables by 60%
 - Improvements based on careful error analysis
 - Simplified square-root implementation



Summer 2015

- Validated ACL2 spec of IEEE 754 Standard against 8M vectors
- Verified all 4 floating-point implementations wrt
 - Sign, exponent, and significand
 - All exceptions
- Found more improvements for future SPARC core
 - Reduced latency further for
 - fdivs, fdivd
 - fsqrts, fsqrt
 - idiv
- Improvements were all proven correct first and then implemented
 - No extensive 24x7 testing needed



Summer 2016

- Applied ACL2 to other areas
 - Temporal properties
 - Proved absence of starvation for certain instructions
- Looked at verification of complex crypto instructions
 - Verified Montgomery algorithms implementations wrt their math specification
 - Saved 25% latency in one instruction
 - Generated test vectors to exercise special cases for these instructions
- Verified and improved fault-tolerant cache-coherency protocol (like CCIX)
 - “Bake off” between ACL2 (TP) and Murphi/PReach (MC)

Summer 2017

- Found more improvements for integer division for future SPARC core
 - Further, significant latency reduction for integer division
 - Again, optimization was first verified, then implemented
- Verified another function using Cellular Automata Shift Registers
- Applied ACL2 to prove absence of starvation for certain instructions
 - Another “bake off” between ACL2 (TP) and SVA model checking

Rockwell Collins Usage



Our Uses of ACL2 to Date

- Microcode Modeling and Proofs
- AAMP7 Information Flow Proofs (GWV Theorem)
 - NSA MILS Accreditation
- Green Hills Information Flow Proofs (GWVr2 Theorem)
 - EAL6+ Accreditation
- AAMP7 Instruction Set Modeling and Proofs
 - Interface to Eclipse-based Debugger
- MicroCryptol Runtime
- Proofs for Guard Prototype (AAMP7 code, vFAAT)
- Data Flow Logic (DFL) for C code
- LLVM Modeling and Proofs
- Other things we can't talk about...

Themes:

- ***Automated High-Level Property Verification for Low-Level Artifacts***
- ***Validation Enabled by Executable Formal Models***

© Copyright 2015 Rockwell Collins, Inc.
All rights reserved.

2

ACL2 Support for Industrial Projects

There have been over 1000 changes to ACL2 since Centaur started using ACL2 in May, 2009. Of those, these were requested by Centaur:

Changes to Existing Features	95
New Features	44
Heuristic and Efficiency Improvements	22
Bug Fixes	72
Changes at the System Level	18
Total due to Centaur	251

Why ACL2 is Successful in Industry

- that was the goal of the project
- efficient, executable logic/programming language with native verifier
- dual-use bit- and cycle-accurate models
- access to Common Lisp programming (via trust tags)
- automatic prover with “a human in the loop”

- encourages development of domain-specific automatic provers allowing *proof maintenance* as designs evolve
- rugged, well documented, free, open source form, many useful books, and a fairly unrestrictive license
- coherent user community devoted to making mechanized verification practical

- industry needs help: their designs are too complicated to get right without mechanized reasoning

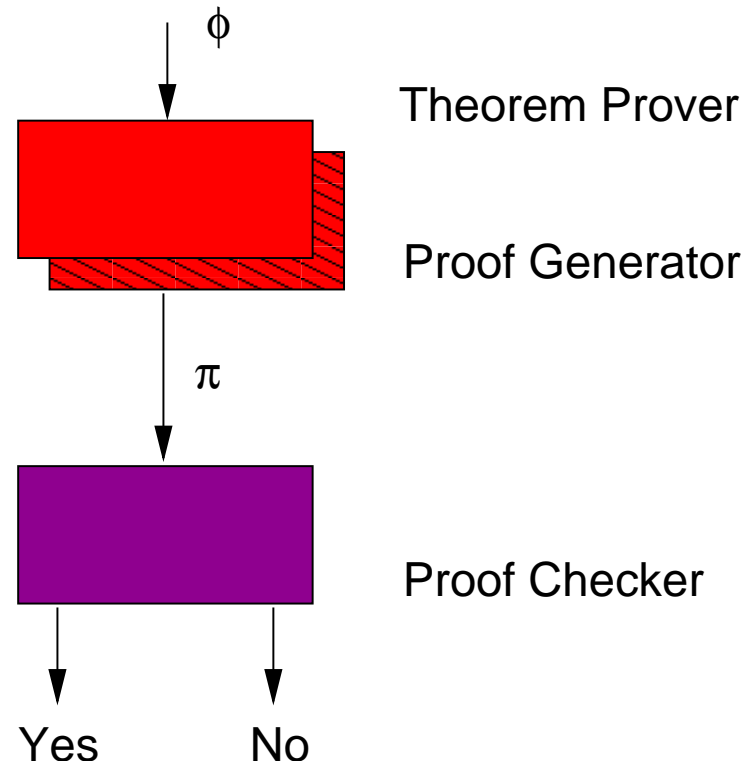
How Do We Know ACL2 is Sound?

“Trust us!” – *Kaufmann and Moore*

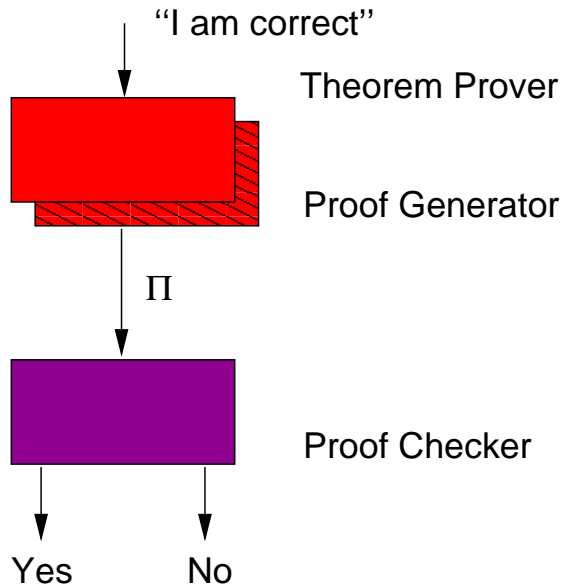
Obviously, we would like to prove it correct.

But with what prover?

Meaning of Correctness



Plan



- Prove “I am correct” with Theorem Prover
- Generate *that* proof Π
- Check Π with Proof Checker
- Never generate another low-level proof

Jared Davis' Stack "Milawa"

