

# The Mechanically Verified Stack Challenge

J Strother Moore  
Inman Chair and Chairman

Department of Computer Sciences  
University of Texas at Austin

# The Premise

The formal methods community is in danger of settling into comfortable niches far smaller than our potential suggests.

To maximize the impact of our science, we must demonstrate that mechanized formal methods can dramatically decrease time-to-market while producing correct systems.

This requires

*pervasive use of formal methods  
throughout the abstraction hierarchy.*

# The Challenge

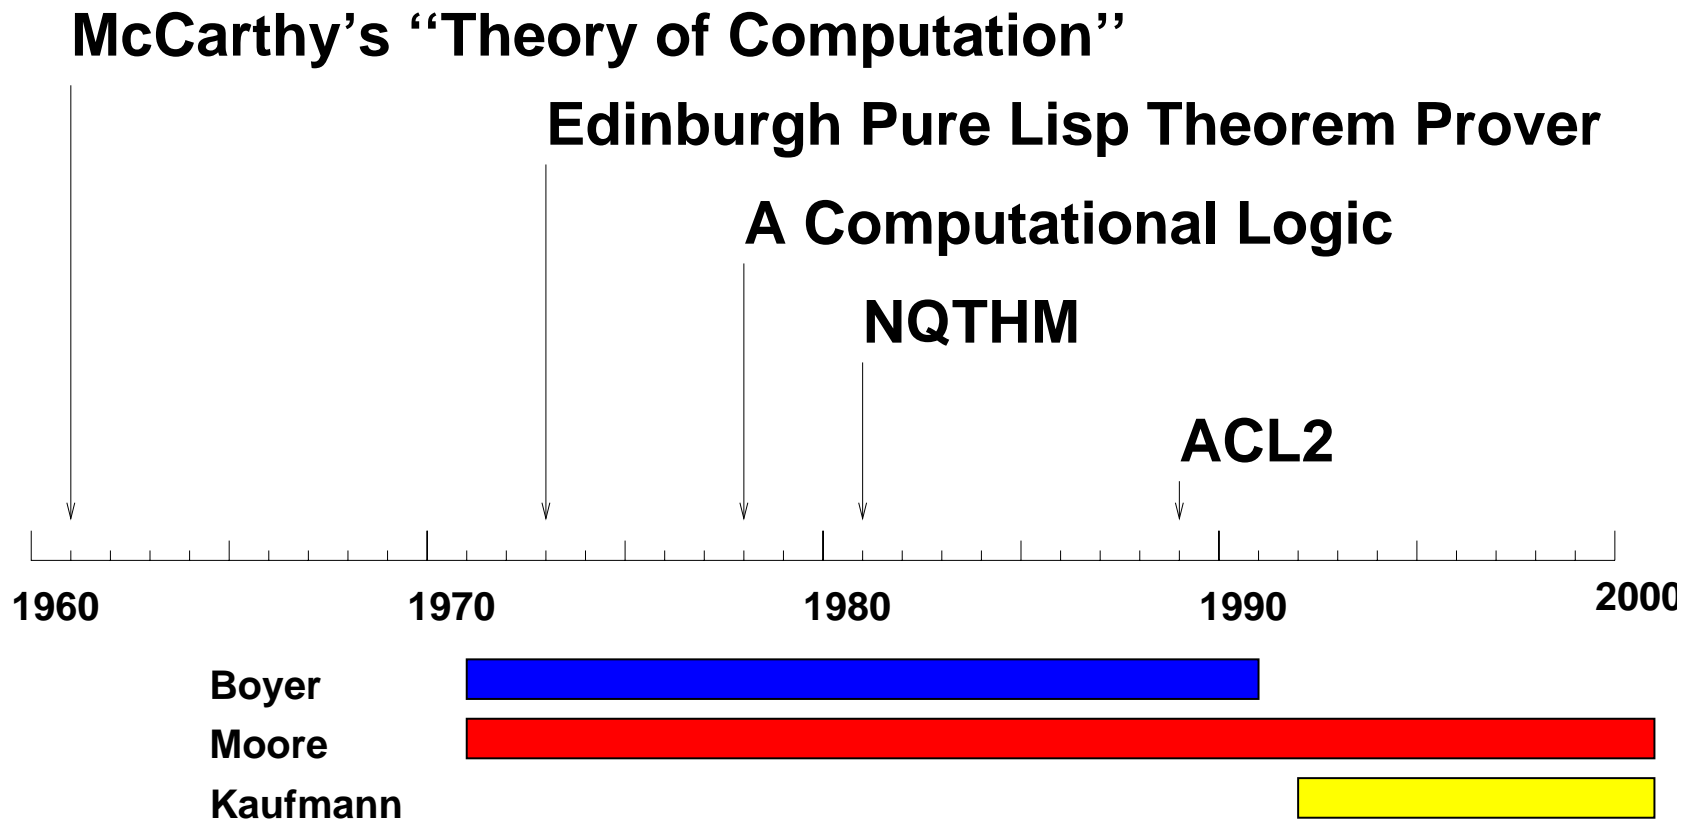
Each major group in the formal methods community should design and mechanically verify a practical embedded system, from transistors to software.

## A Personal Perspective

Instead of debugging a program, one should prove that it meets its specifications, and this proof should be checked by a computer program.

— *John McCarthy, “A Basis for a Mathematical Theory of Computation,” 1961*

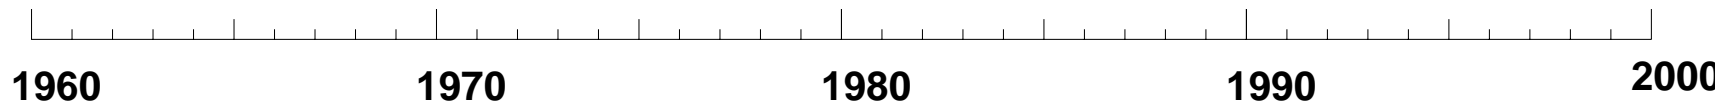
# Boyer-Moore Project



**simple list processing**

**academic math and cs**

**commercial  
applications**



# 1970s

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

```
(theorem
  (equal (append (append a b) c)
         (append a (append b c))))
```



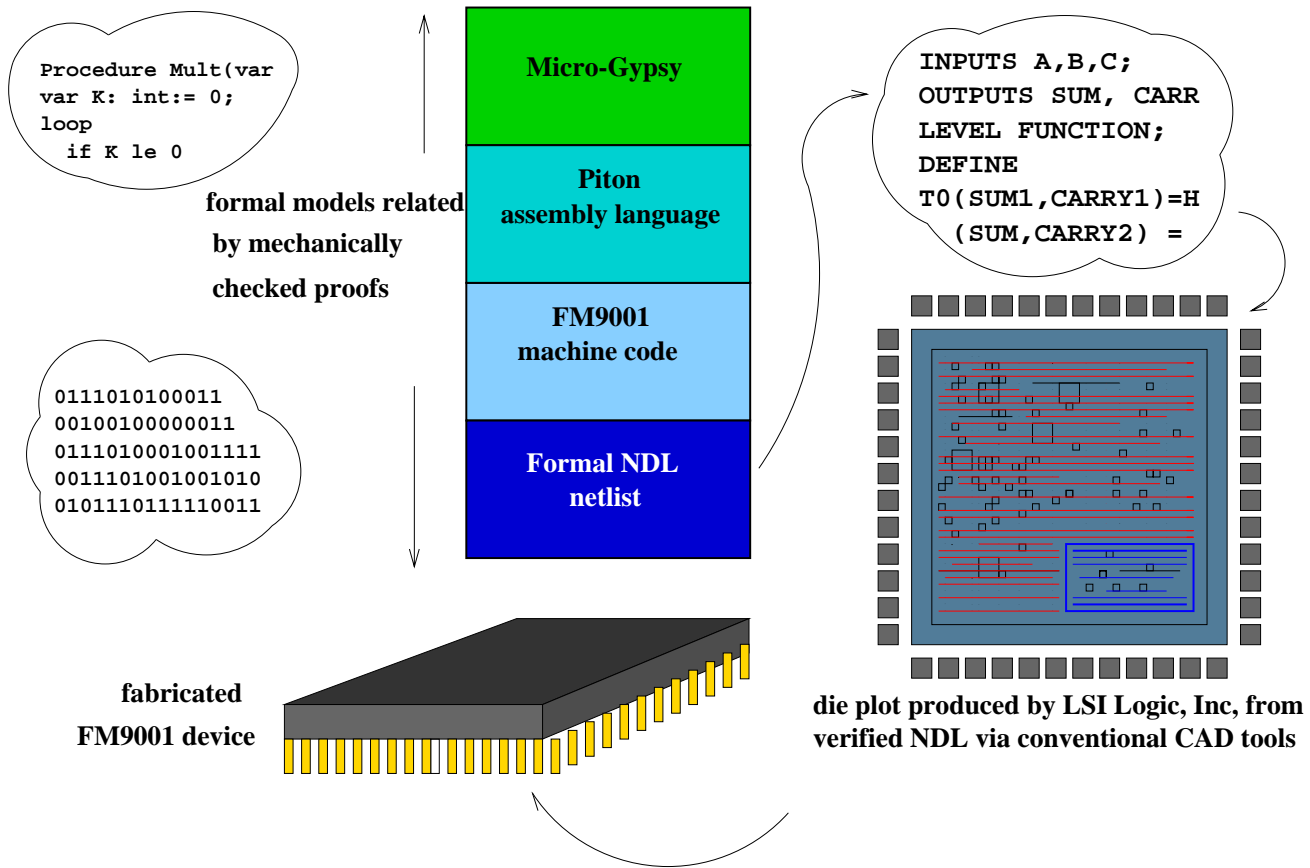
# 1980s Academic Math

- invertibility of RSA encryption
- undecidability of the halting problem
- Gödel's First Incompleteness Theorem (Shankar)
- Gauss' law of quadratic reciprocity (Russinoff)

# 1980s Academic CS

- microprocessor: gates to machine code (Hunt)
- assembler-linker-loader (Moore and Kaufmann)
- compiler (Young)

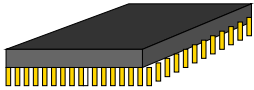
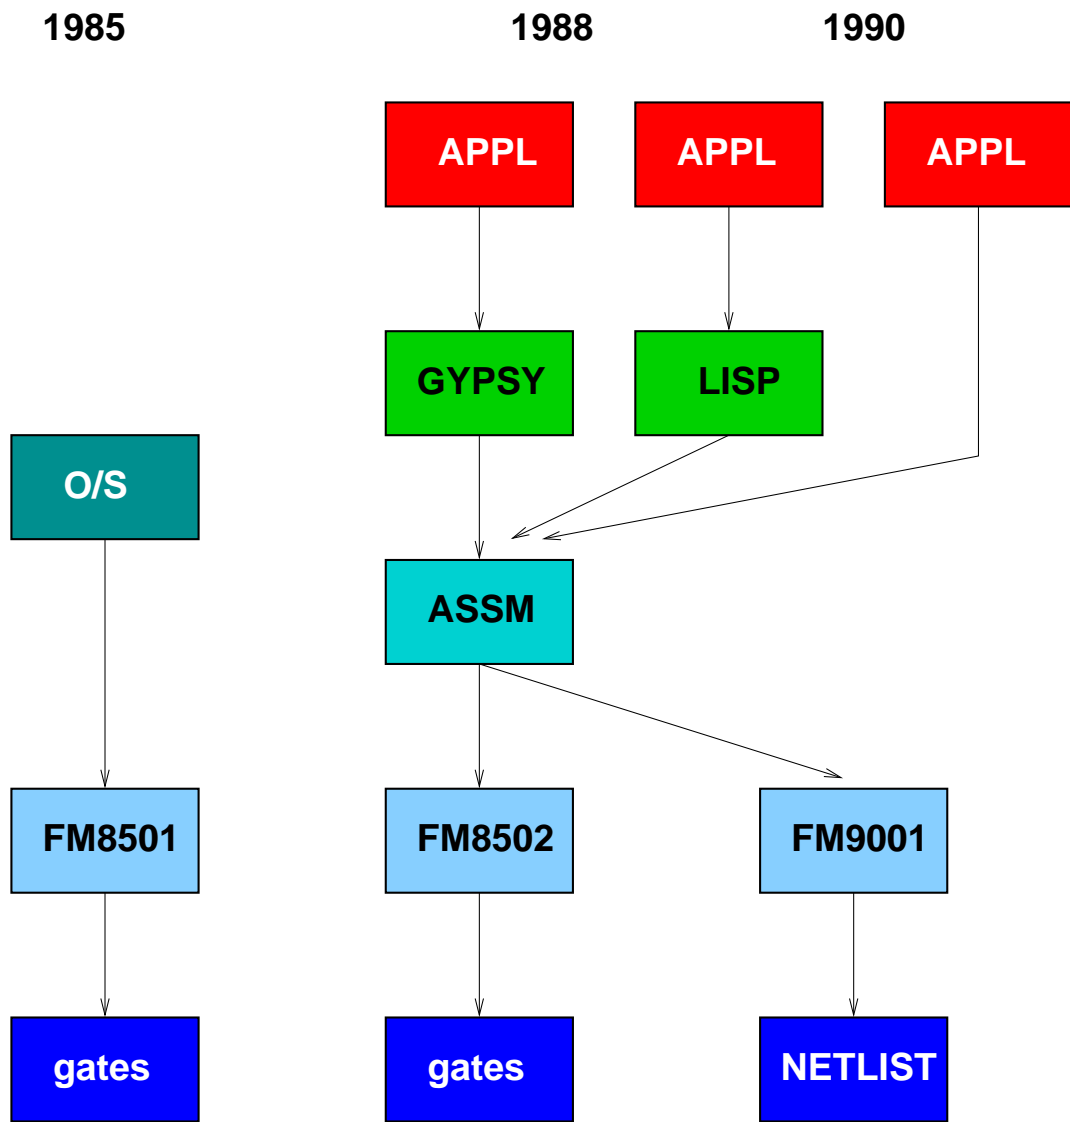
# The CLI Stack



# Additional Stack Components

- operating system (Bevier)
- Lisp compiler (Flatau)
- applications (Wilding)

- Applications were mechanically verified with respect to the high-level languages.
- Correct binary images were produced by mechanically verified tools.
  - `compile`
  - `assemble`
  - `link`
  - `load`



# What's Missing?

- The FM9001 was modeled at the gate level.
- The FM9001 had an unrealistically simple architecture: memory-mapped io, no pipeline, no speculation, no floating point, crude interrupts, no cache.

- None of the programming languages had io or floating point.
- The high-level languages were too simple to be of practical use.
  - The micro-Gypsy language had very few primitives and no dynamically allocated data types (e.g., records).
  - The Pure Lisp had automatic storage



allocation (e.g., `cons`) but no verified garbage collector.

- The compilers, assembler, linker, and loader were “cross-platform” transformers.
- No useful application programs were verified.

- The operating system was not hosted on the FM9001.

## **What Was Learned?**

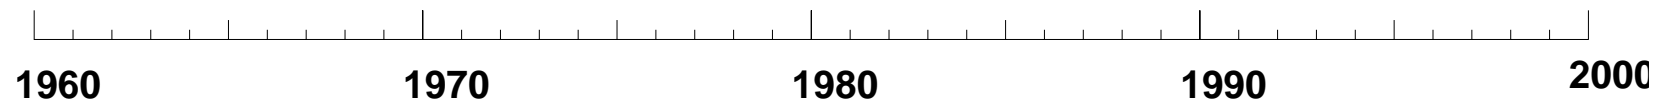
System verification does not require super-human skills.

It requires better tools than we had in 1985.

**simple list processing**

**academic math and cs**

**commercial  
applications**



**stack project**



**ACL2 development**

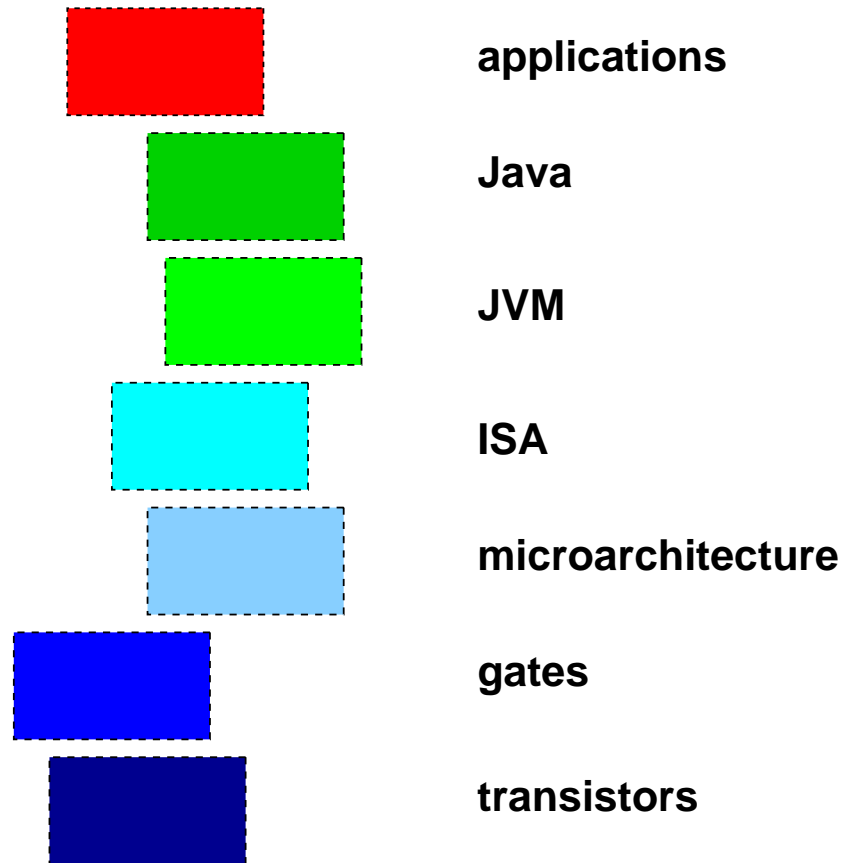
# 1990s

- Motorola 68020 and Berkeley C String Library
- Motorola CAP DSP
- AMD K5 FDIV
- AMD Athlon fp

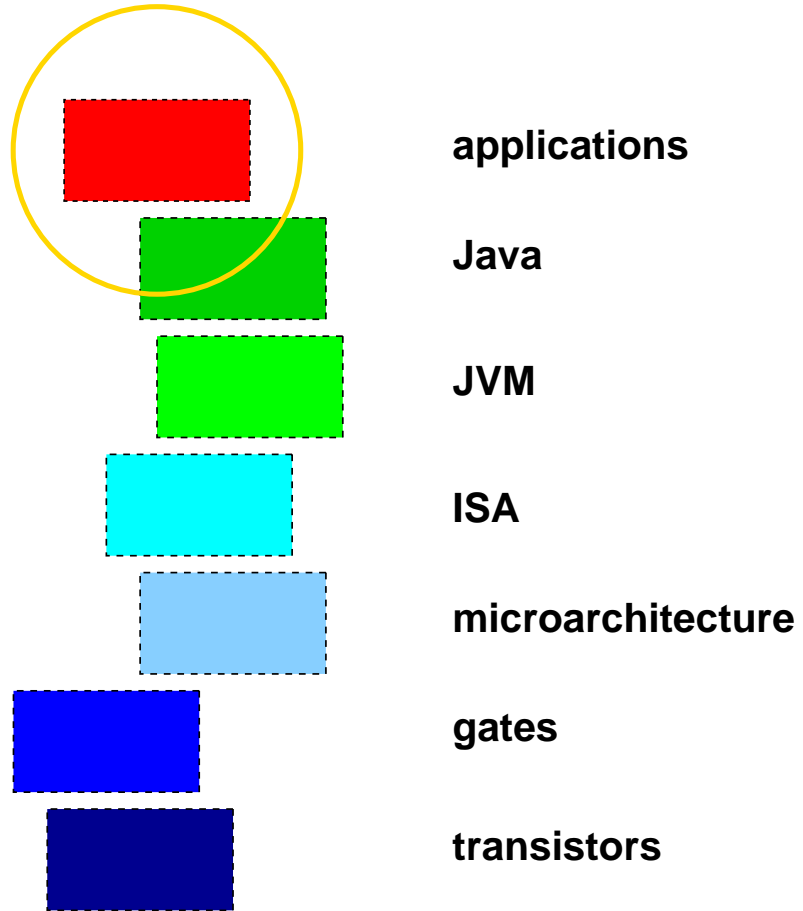
- Rockwell Collins JEM1
- IBM Cryptographic coprocessor
- Union Switch & Signal and Argonne Nat'l Labs checkers
- FM9801
- JVM

- Java classes
- IBM Array Verification Tool
- Rockwell Collins CAPS

# An Imaginary Stack



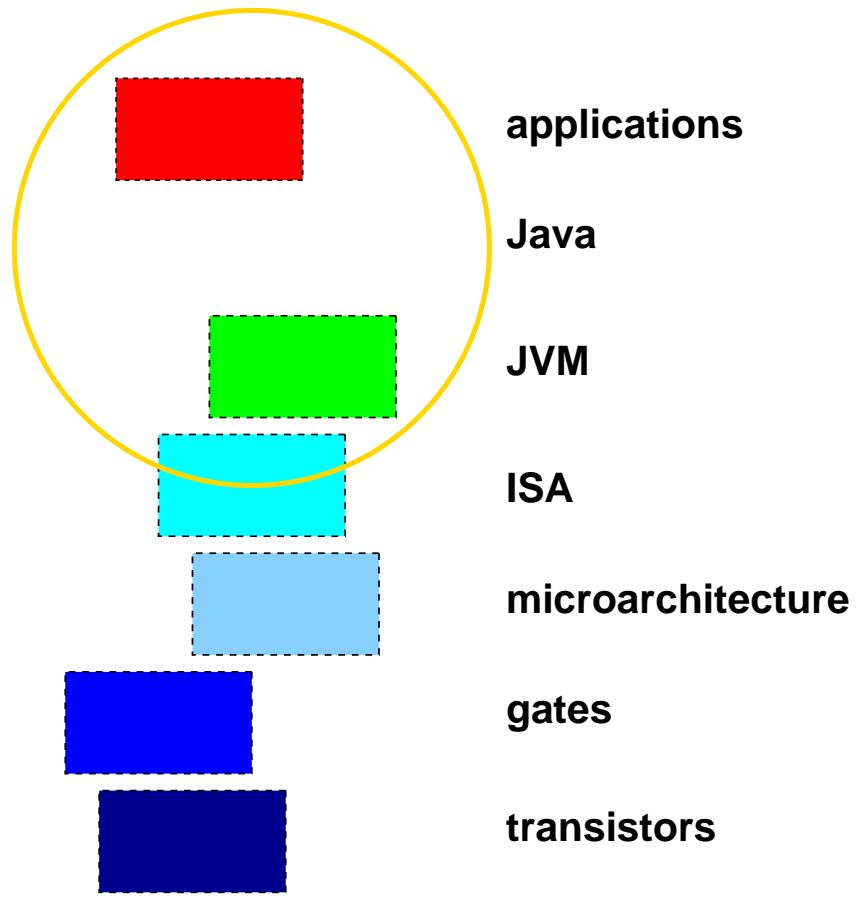




```
class Container {
    public int counter; }

class Job extends Thread {
    Container objref;
    public Job incr () {
        synchronized(objref) {
            objref.counter = objref.counter + 1; }
        return this; }
    public void setref(Container o) {
        objref = o; }
    public void run() {
        for (;;) {incr(); } } }
```

```
class Apprentice {
    public static void main(String[] args){
        Container container = new Container();
        for (;;) {Job job = new Job();
                job.setref(container);
                job.start(); } } }
```



**applications**

**Java**

**JVM**

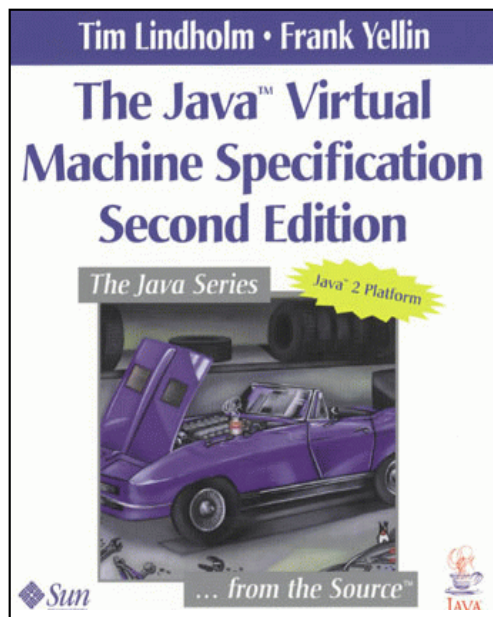
**ISA**

**microarchitecture**

**gates**

**transistors**

# Java (Golden, Krug, Liu, Moore, Porter)

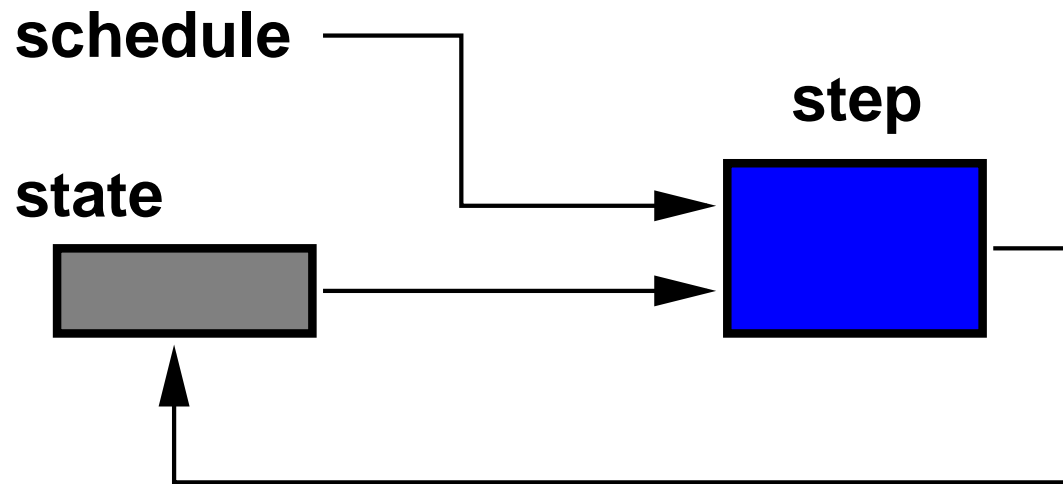


```
; JVM in ACL2
; J Moore and George Porter

(defun make-state (tt hp ct)
  ...)

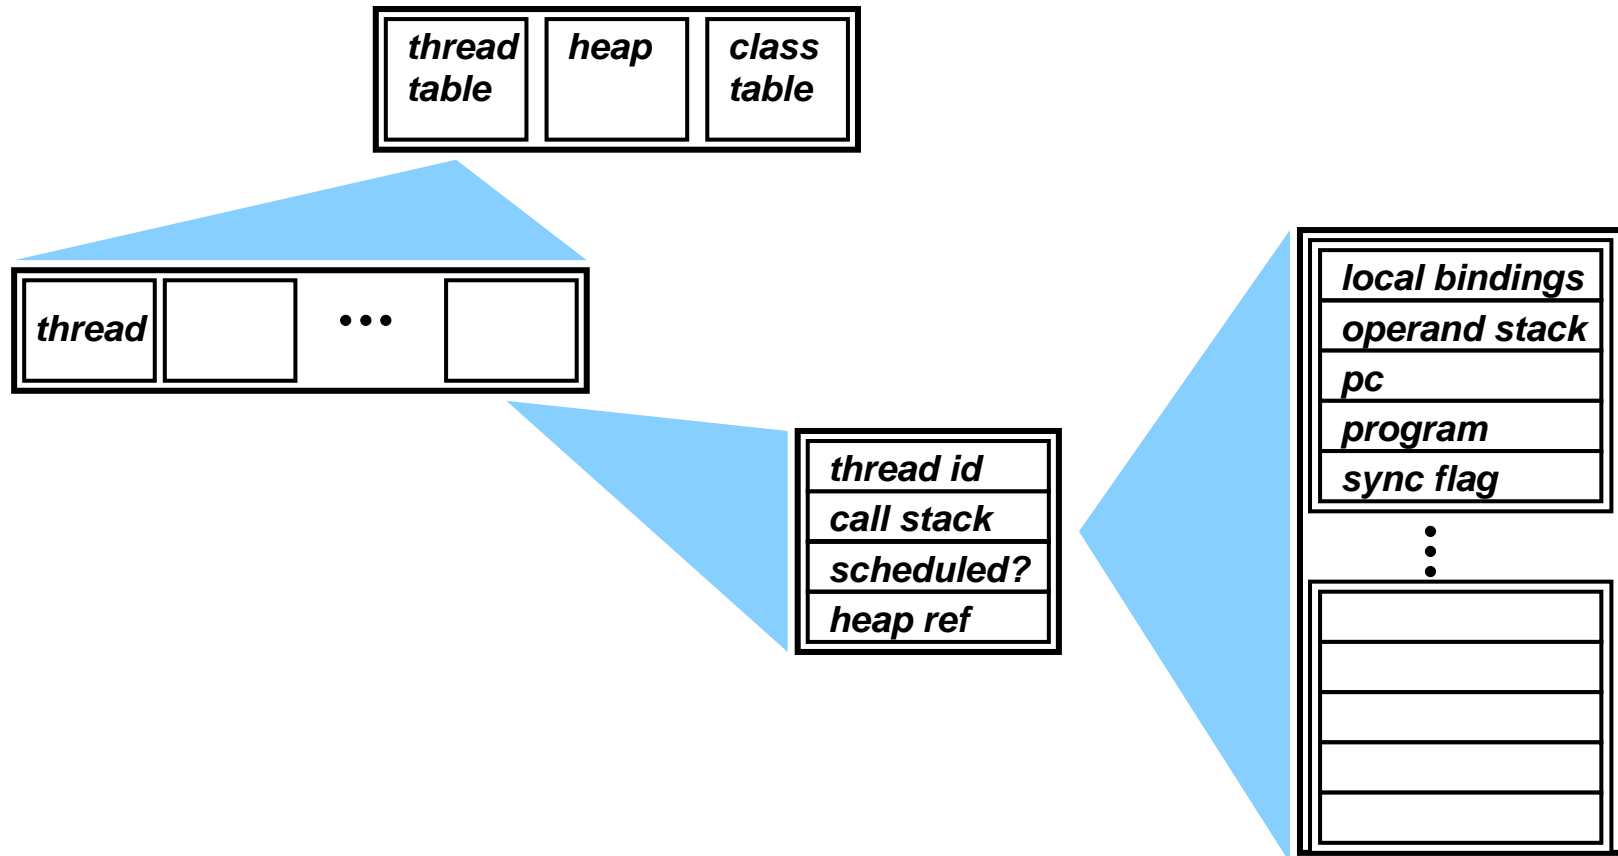
(defun step (th s)
  ...)

(defun run (sched s)
  (if (endp sched)
      s
      (run
       (cdr sched)
       (step (car sched) s))))
```



```
(defun run (schedule state)
  (if (endp schedule)
      state
      (run (cdr schedule)
           (step (car schedule) state))))
```

# Our State: $\langle tt, hp, ct \rangle$



```
(defun step (th s)
  (if (equal (call-stack-status th s)
            'SCHEDULED)
      (do-inst (next-inst th s) th s)
      s))
```



```
(defun do-inst (inst th s)
  (case (op-code inst)
    (AALOAD (execute-AALOAD inst th s))
    (AASTORE (execute-AASTORE inst th s))
    (ALOAD (execute-ALOAD inst th s))
    (ALOAD_0 (execute-ALOAD_X inst th s 0))
    (ALOAD_1 (execute-ALOAD_X inst th s 1))
    (ALOAD_2 (execute-ALOAD_X inst th s 2))
    (ALOAD_3 (execute-ALOAD_X inst th s 3))
    ...))
```

```

(defun execute-AALOAD (inst th s)
  (let* ((frame (top-frame th s))
         (index (top (stack frame)))
         (aref (top (pop (stack frame))))
         (array (deref aref (heap s))))
    (modify th s
      :pc (+ (inst-length inst) (pc frame))
      :stack
        (push (element-at index array)
              (pop (pop (stack frame)))))))

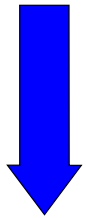
```

## Some Java/JVM/ACL2 Code

```
public static int fact(int n){  
  if (n<=0) return 1;  
  else return n*fact(n-1);}
```

javac	jvm2acl2
Method int fact(int)	("fact" (INT) nil
0 iload_1	(ILOAD_1)
1 ifgt 5	(IFGT 5)
4 iconst_1	(ICONST_1)

**.java**



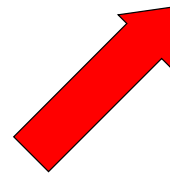
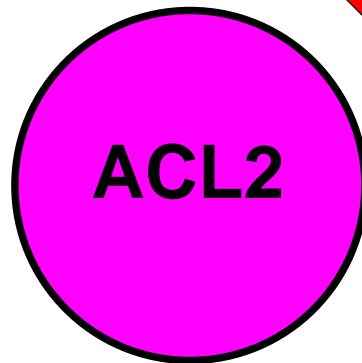
**javac**

**.class**



**jvm2acl2**

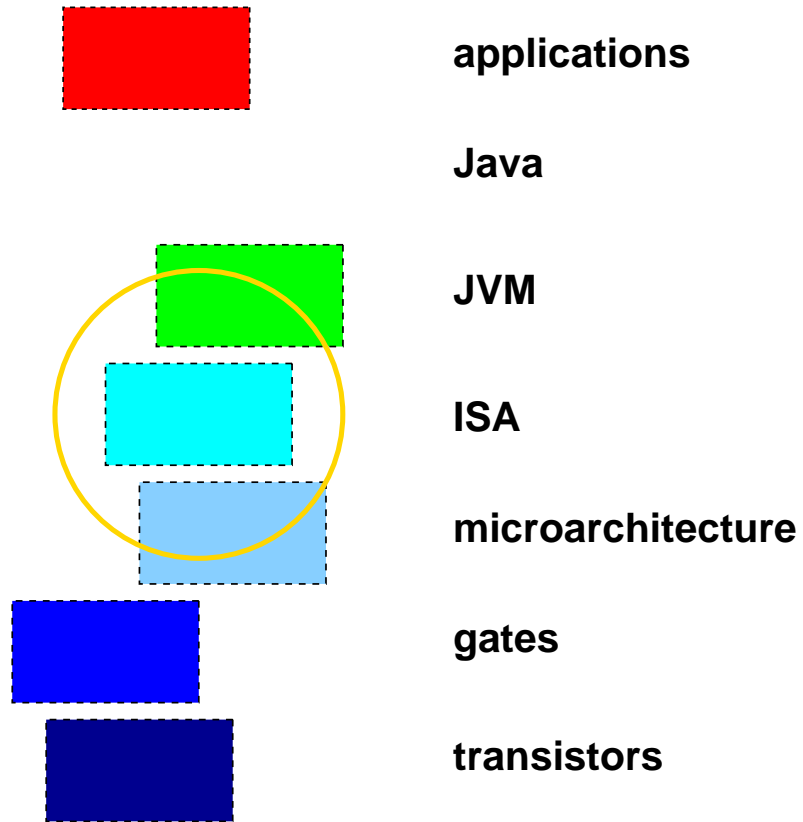
**.lisp**



**Theorems**

**“fact(5)=120”**

**“fact(n)=n!”**



# Rockwell-Collins / aJile Systems JEM1

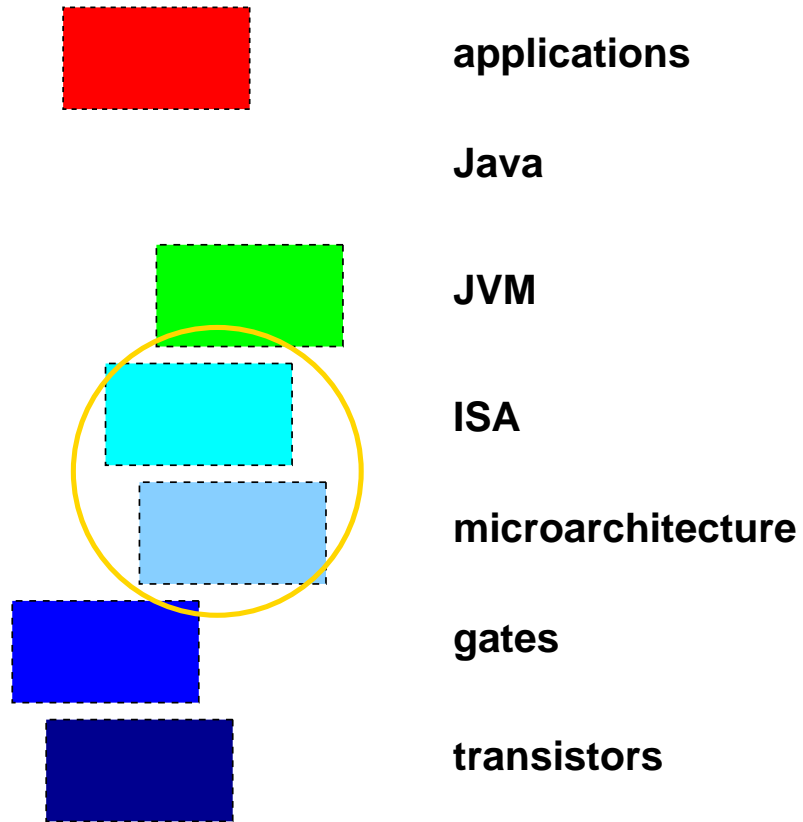


The world's first silicon Java Virtual Machine was first modeled in ACL2.

The formal model is executable.

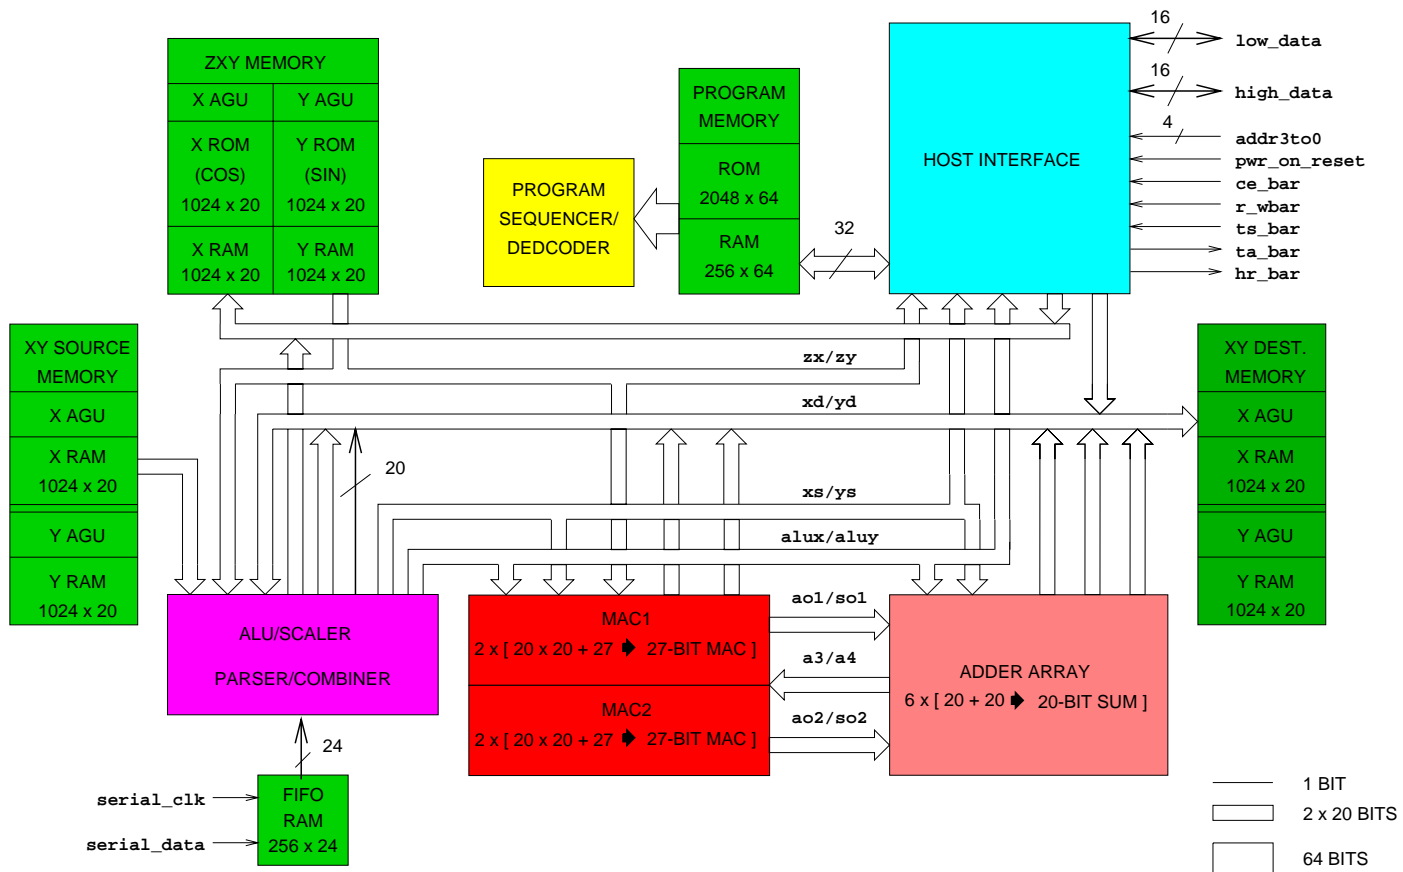
It was used in place of a C simulator for requirements and certification testing.

It runs at 90% the speed of the C simulator.





# Motorola CAP DSP (Brock)





**ROM containing  
50 microcoded  
DSP algorithms**

**Pipelined  
microarchitecture**

**Sequential  
microcode ISA**

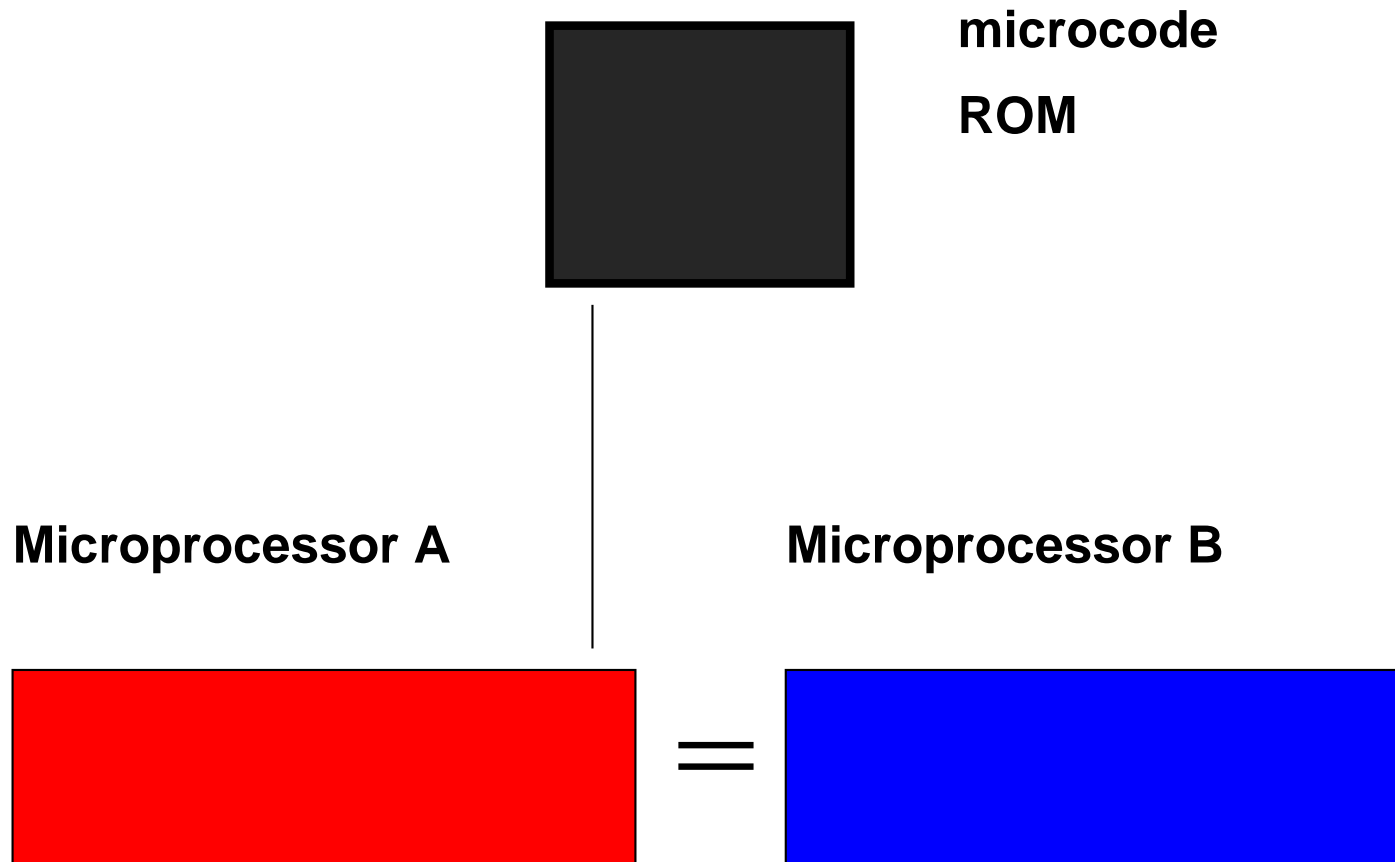


**=**



**(if no hazards)**

# Rockwell Collins Avionics



## 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.  $fdiv = qq_1 + q_0$  *mode*

# 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:

$$1/12 = 0.08\overline{33} \approx 0.083 = sd_2$$

Quotient Digit Calculation:

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

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

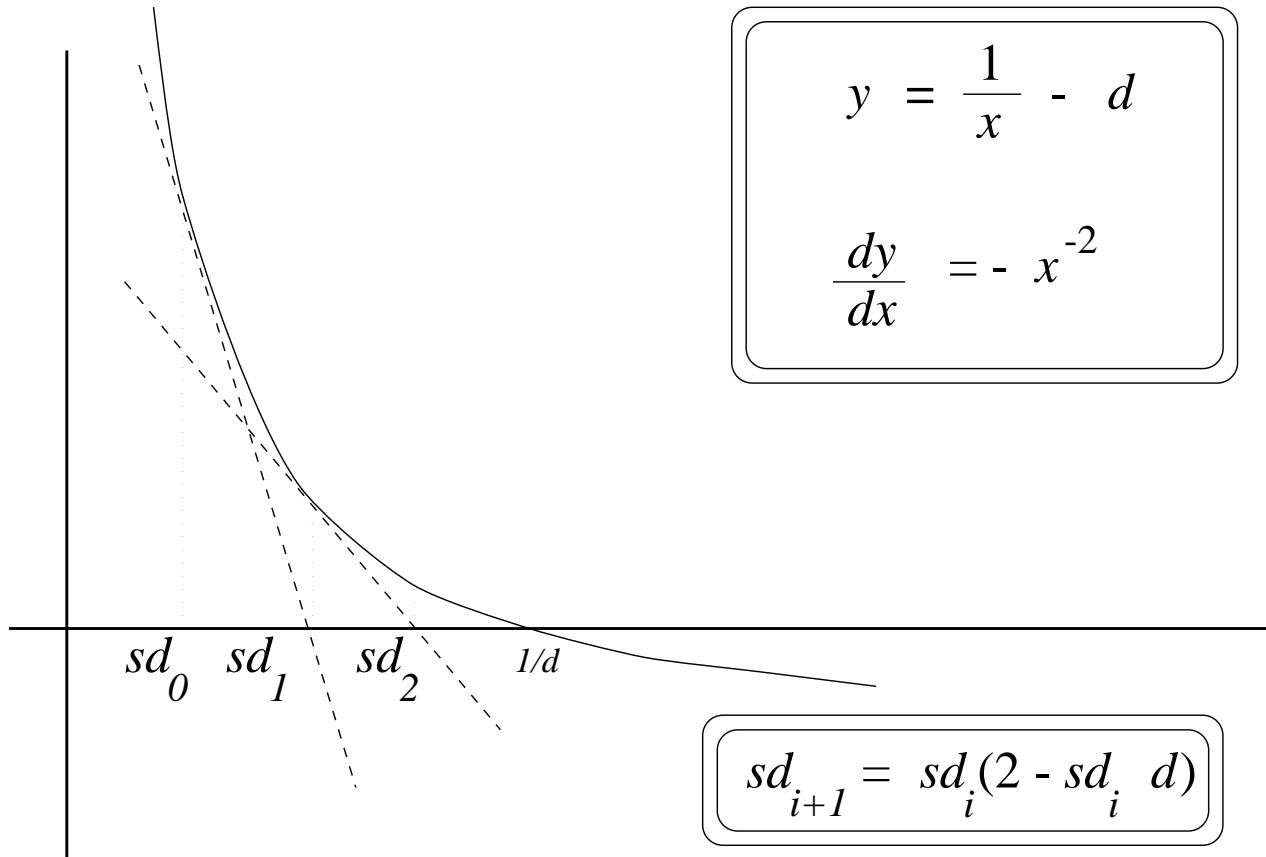
$$0.083 \times .0400 = .0033200 \approx .003400 = q_2$$

$$0.083 \times -.0008 = -.0000664 \approx -.000067 = q_3$$

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

# The Formal Model of the Code

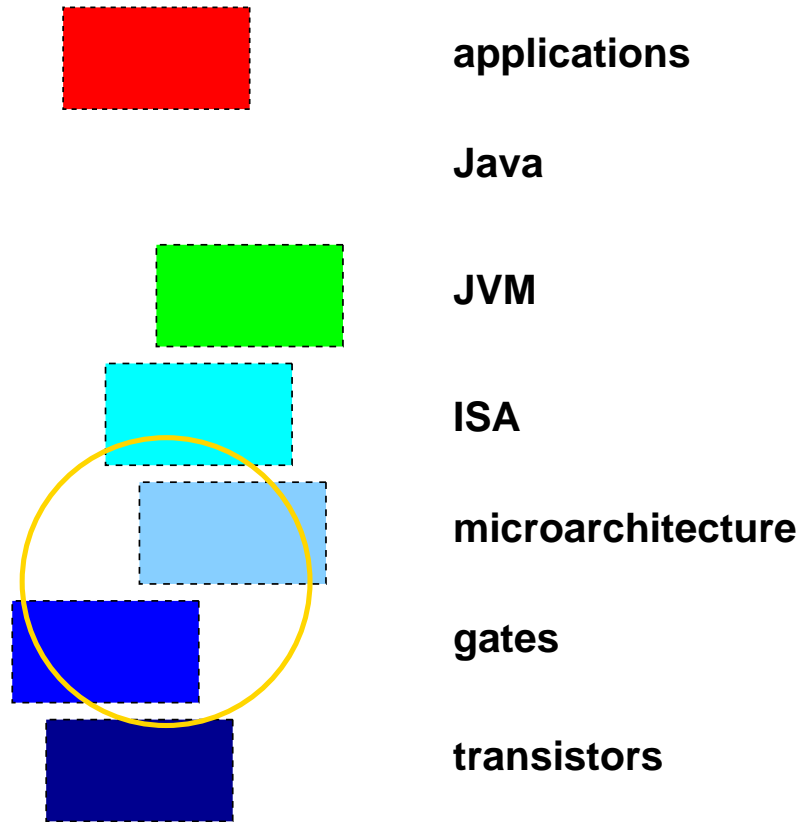
```
(defun FDIV (p d mode)
  (let*
    ((sd0 (eround (lookup d)                '(exact 17 8)))
     (dr  (eround d                          '(away 17 32)))
     (sdd0 (eround (* sd0 dr)                '(away 17 32)))
     (sd1  (eround (* sd0 (comp sdd0 32))    '(trunc 17 32)))
     (sdd1 (eround (* 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 ... fddiv)
        fddiv)))
```



# The K5 FDIV Theorem

```
(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))))
```

(by Moore, Lynch and Kaufmann, in 1995,  
*before the K5 was fabricated*)





AMD ATHLON™  
ROCKETS TO **1GHz**



All elementary floating-point operations on the **AMD Athlon** were

- specified in ACL2 to be IEEE compliant,
- proved to meet their specifications, and
- the proofs were checked mechanically – before fab.

(Russinoff and Flatau)

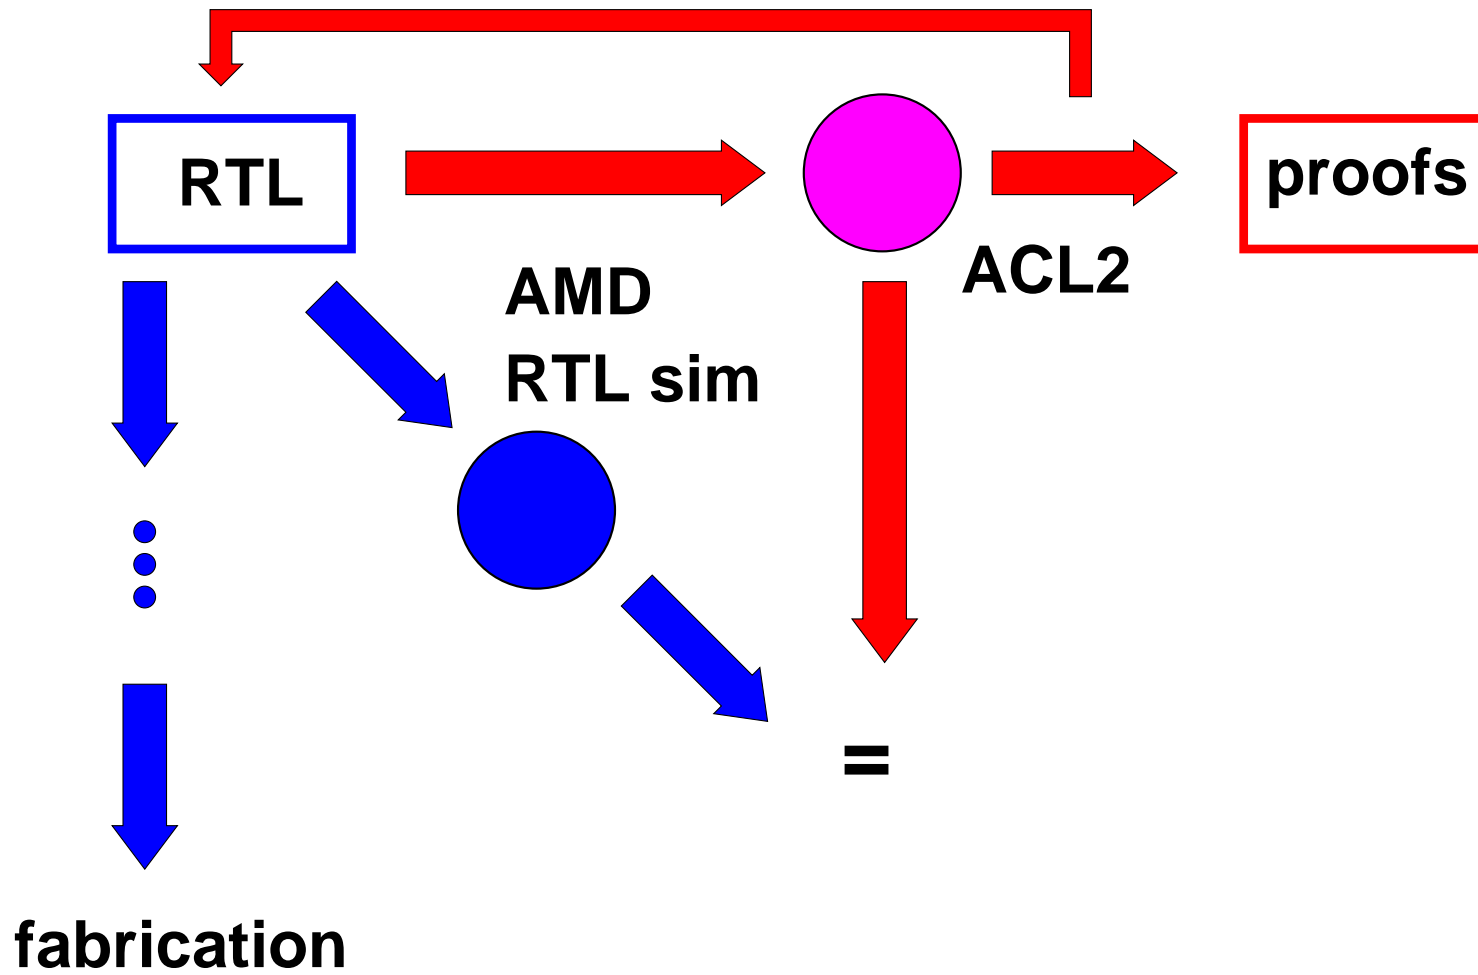
# AMD Athlon FMUL

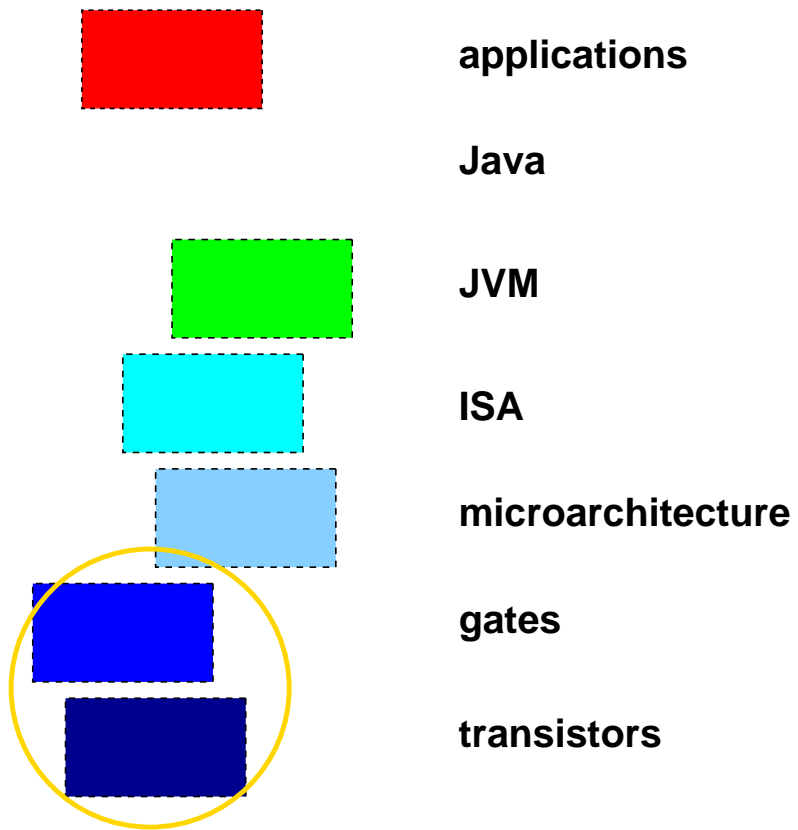
```
module FMUL; // sanitized from AMD Athlon(TM)
            // by David Russinoff and Art Flatau
//*****
// Declarations
//*****
//Precision and rounding control:
`define SNG    1'b0    // single precision
`define DBL    1'b1    // double precision
`define NRE    2'b00   // round to nearest
`define NEG    2'b01   // round to minus infinity
`define POS    2'b10   // round to plus infinity
```

```

//Parameters:
input x[79:0];           //first operand
input y[79:0];           //second operand
input rc[1:0];           //rounding control
input pc;                //precision control
output z[79:0];          //rounded product
//*****
// First Cycle
//*****
//Operand fields:
sgnx = x[79]; sgnx = y[79];
expx[14:0] = x[78:64]; expy[14:0] = y[78:64];

```





**applications**

**Java**

**JVM**

**ISA**

**microarchitecture**

**gates**

**transistors**



# Summary of Our Current State

We have mechanically verified

- RTL for industrial designs,
- commercial microcode, and
- simple programs in a widely used programming language.

In each industrial site where ACL2 is being used, proprietary formal information has been created:

- models of artifacts,
- definition of “in-house” concepts,
- specifications, and
- lemmas and proof strategies.

Our lowest level formal models are usually

- “bit accurate,”
- “cycle accurate,” and
- efficiently executable (comparable to C).

Our formal models often replace conventional simulation models:

- they are as accurate and
- they run about as fast, but
- they can be formally analyzed.

Formal models add value.

## The evidence supports

- the value-added proposition,
- the claim that formal methods requires less manpower than testing, and
- the claim that formal methods reduces time-to-market.

We have certified books (lemma libraries) for

- arithmetic and bit-vectors,
- stuttering bi-simulation,
- a BDD package (60% CUDD speed), and
- a sound and complete mu-calculus model checker.

# The Challenge

Each major group in the formal methods community should design and mechanically verify a practical embedded system, from transistors to software.

# Why Build a Stack?

- There is still much to be learned.
- Practice makes perfect (more accurately, practice encourages mechanization).
- A stack forces you to deliver what you assume.



# The Rules of the Game

- The project should produce an artifact, e.g., a chip.
- The artifact's behavior should be of interest to people not in formal methods.
- The artifact should come with a “warranty” expressed as a theorem.

- The warranty should be certified mechanically; user input and interaction are allowed but a machine must be responsible for the soundness of the claim.
- The machine used to certify the warranty should be available for others, at least, to use and test, if not inspect.

## Why It is Hard

There are unsolved technical problems.

But mainly:

None of our systems allow us to move up and down the abstraction hierarchy with the ease required to build a practical system.

# The Challenge Will Encourage

- more automation,
- integration of formal approaches,
- development of collaborative environments,

- identification of assumptions,
- specification of crucial system interfaces,  
and
- elegant designs of practical systems.