

A Career in Research: Mike Gordon and Hardware Verification

Lawrence C Paulson

Automated Reasoning Workshop 2018

The world of computing in 1975

16MB a *colossal*
amount of memory



mainframe disk, about 5MB



minicomputers,
e.g. 16KB memory



The world of theory, 1975–80

C. A. R. Hoare. An axiomatic basis for computer programming
Communications of the ACM **12** (10), Oct. 1969

Dana Scott. Outline of a mathematical theory of computation.
Technical Report PRG-2, University of Oxford, Nov. 1970.

operational semantics
just emerging

denotational semantics and
fixed-point theory

R. Milner. *A Calculus of Communicating Systems*. Springer, 1980.

type theory emerging

process algebras

Edinburgh LCF (1975–1979)

M. J. C. Gordon, R. Milner, and C. P. Wadsworth.
Edinburgh LCF: A Mechanised Logic of Computation.
LNCS 78. Springer, 1979.

- ❖ The first real *proof assistant* (for computation theory)
- ❖ Introducing ML (the first *polymorphic* functional language)

With so many nascent fields,
what did Mike decide to do?

Software is being solved, so let's verify *hardware*.

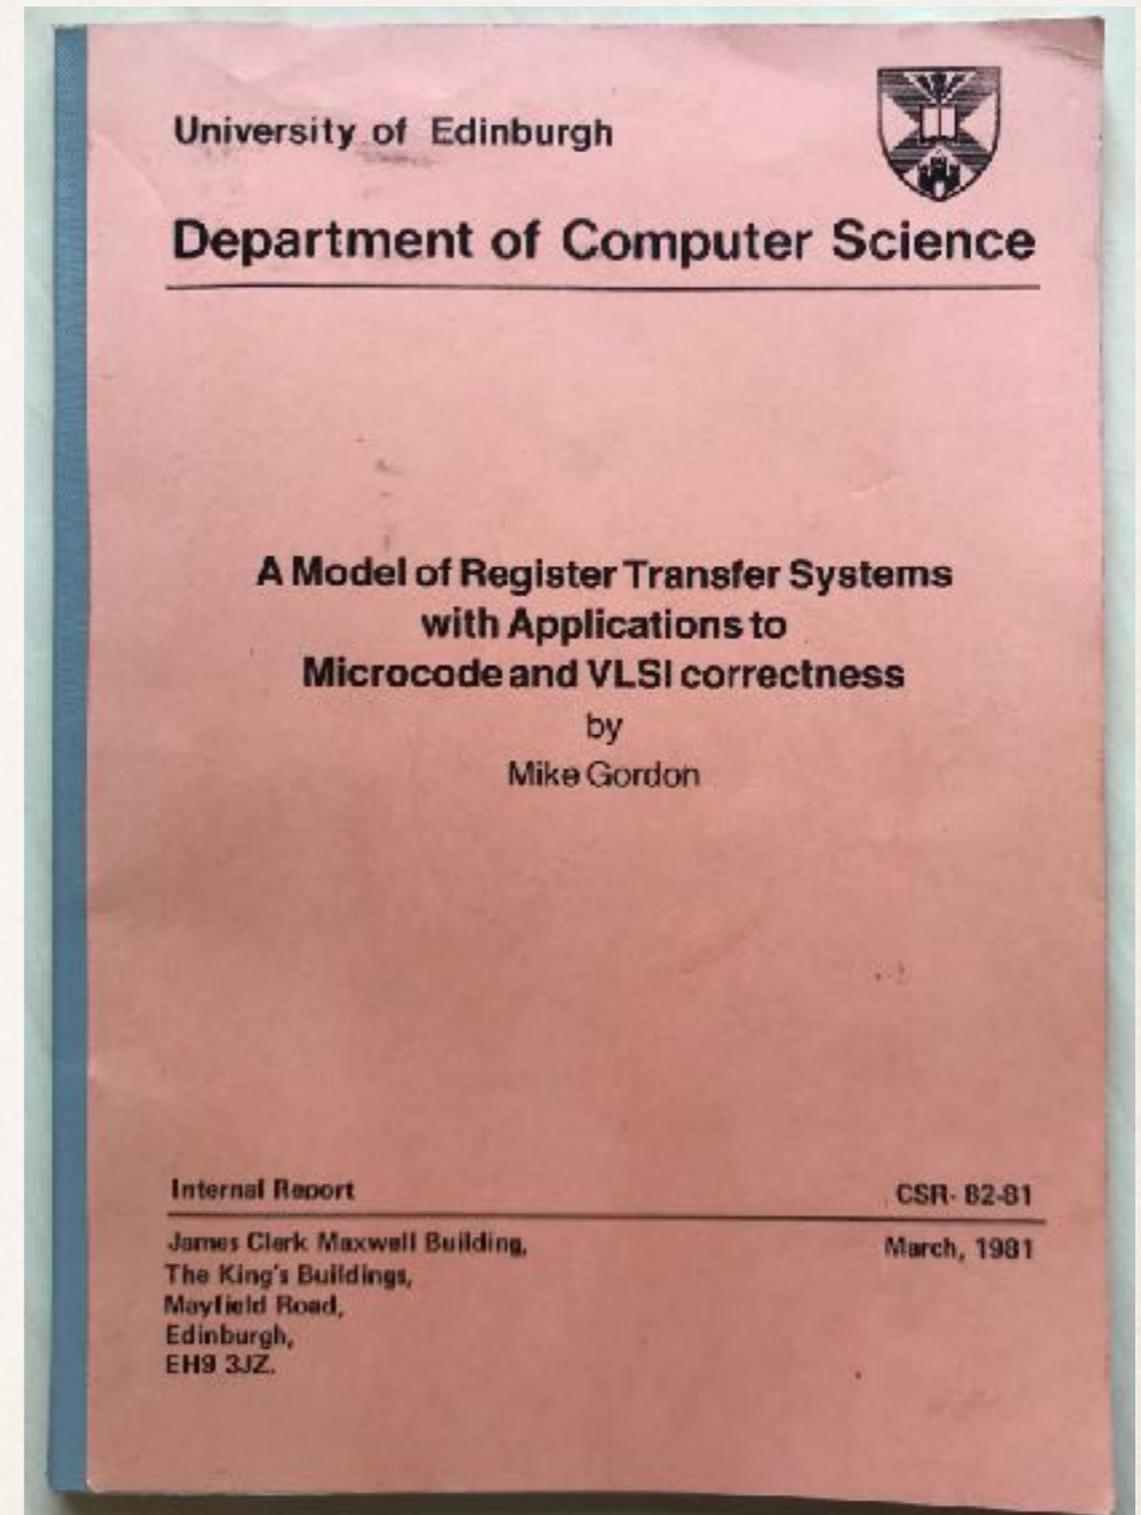
He talked to the hardware experts at Edinburgh

and sketched out some theory

... and designed his own computer!

A model of register transfer systems
with applications to microcode and
VLSI correctness.

Technical Report CSR-82-81,
University of Edinburgh, Mar. 1981.



the circuitry

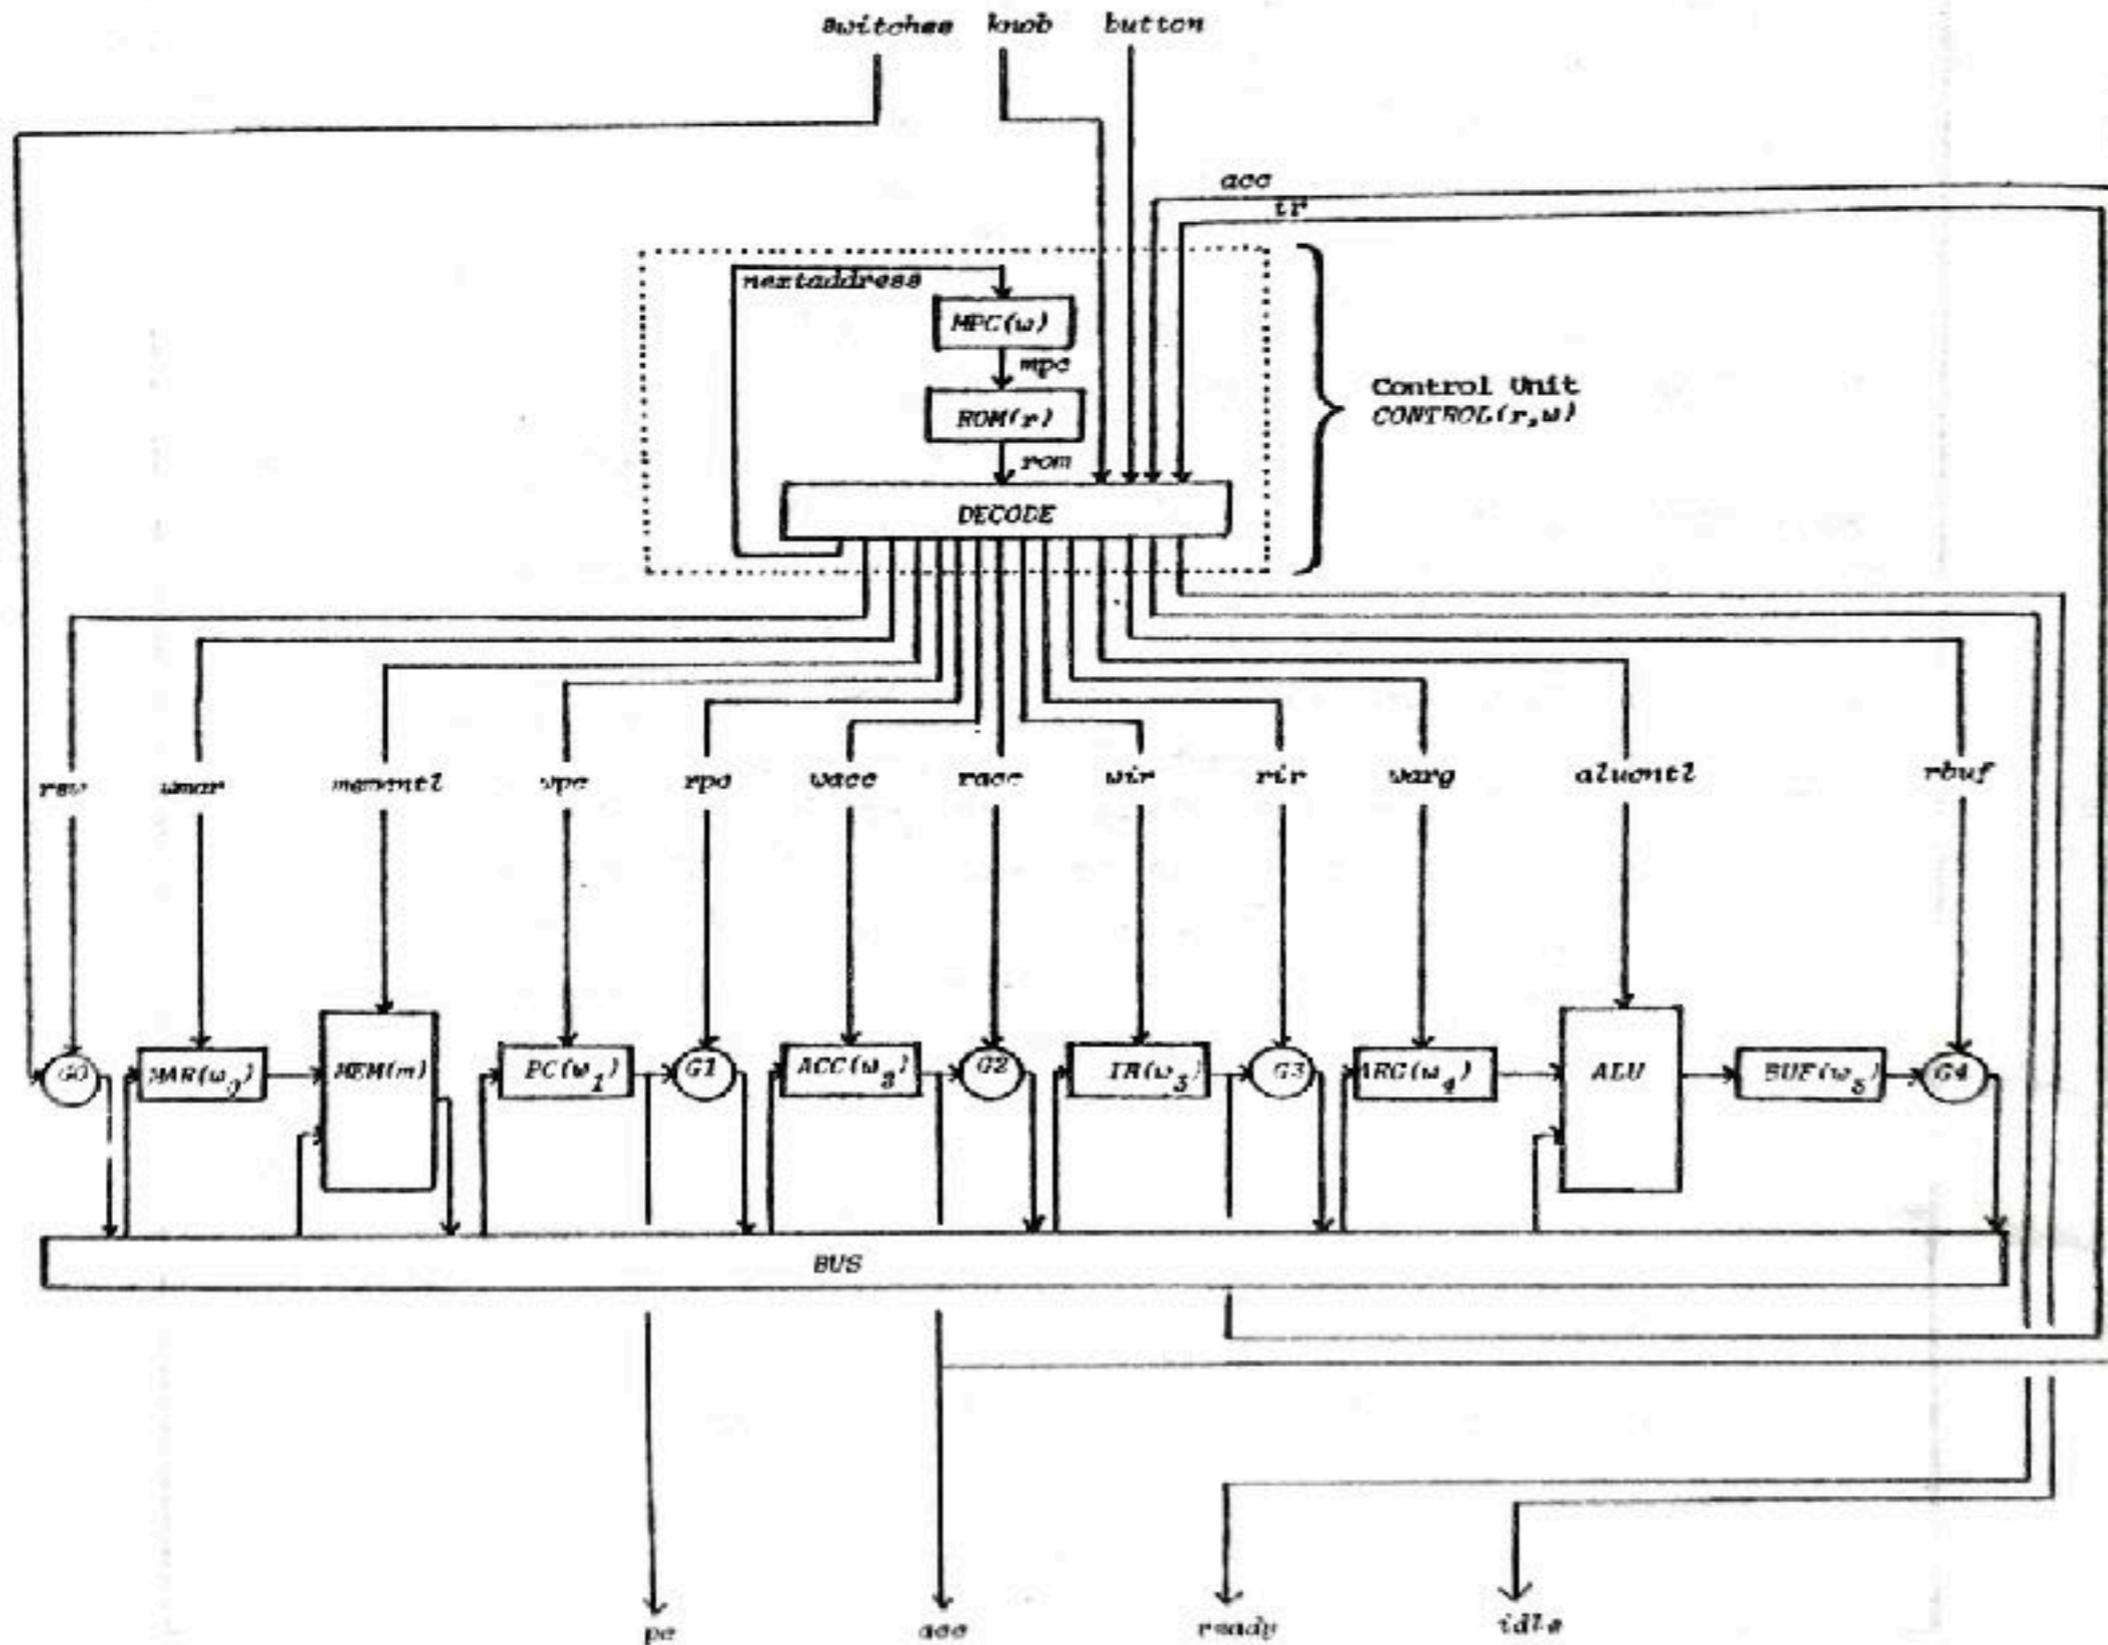


Fig. XIV: The host machine for implementing the computer

a microprogram

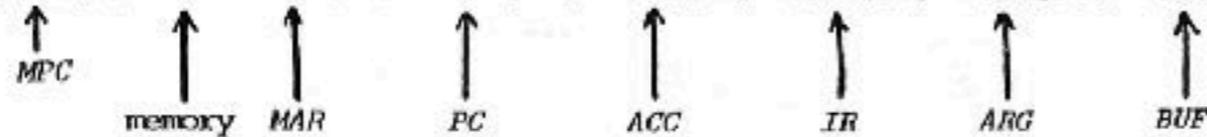
<u>ROM</u>	<u>Microinstruction</u>	<u>Comments</u>
<u>address</u>		
0	<i>ready, idle; button + 1, 0</i>	branch to 1 if button pressed, otherwise loop
1	<i>; knob + 1</i>	decode knob position
2	<i>rsu, wpc ; 0</i>	<i>switches + PC</i>
3	<i>rsu, wacc ; 0</i>	<i>switches + ACC</i>
4	<i>rpe, wmar ; 7</i>	<i>PC + MAR</i>
5	<i>ready ; button + 0, 6</i>	begin fetch-decode-execute cycle
6	<i>rpe, wmar ; 8</i>	<i>PC + MAR</i>
7	<i>racc, write; 0</i>	<i>ACC + MEM(MAR)</i>
8	<i>read, wir ; 9</i>	<i>MEM(MAR) + IR</i>
9	<i>; opcode + 10</i>	decode
10	<i>; 0</i>	halt
11	<i>rir, wpc ; 6</i>	<i>JMP: IR + PC</i>
12	<i>; acc=0 + 11, 17</i>	<i>JZRO:</i>
13	<i>rdoc, warg ; 19</i>	<i>ADD: ACC + ARG</i>
14	<i>racc, warg ; 22</i>	<i>SUB: ACC - ARG</i>
15	<i>rir, wmar ; 24</i>	<i>LD: IR → MAR</i>
16	<i>rir, wmar ; 25</i>	<i>ST: IR → MAR</i>
17	<i>rpe, inc ; 18</i>	<i>PC + 1 + BUF</i>
18	<i>rbuf, wpc ; 5</i>	<i>BUF + PC</i>
19	<i>rir, wmar ; 20</i>	<i>IR + MAR</i>
20	<i>read, add ; 21</i>	<i>ARG + MEM(MAR) + BUF</i>
21	<i>rbuf, wacc ; 17</i>	<i>BUF → ACC</i>
22	<i>rir, wmar ; 23</i>	<i>IR + MAR</i>
23	<i>read, sub ; 21</i>	<i>ARG - MEM(MAR) + BUF</i>
24	<i>read, wacc ; 17</i>	<i>MEM(MAR) → ACC</i>
25	<i>racc, write; 17</i>	<i>ACC + MEM(MAR)</i>

Fig. XVI Microinstruction in control units' ROM: the microprogram microcode

internal specifications

HOSTMACHINE = (S, out, next)

S = Word[5] × Mem × Word[13] × Word[13] × Word[16] × Word[16] × Word[16] × Word[16]



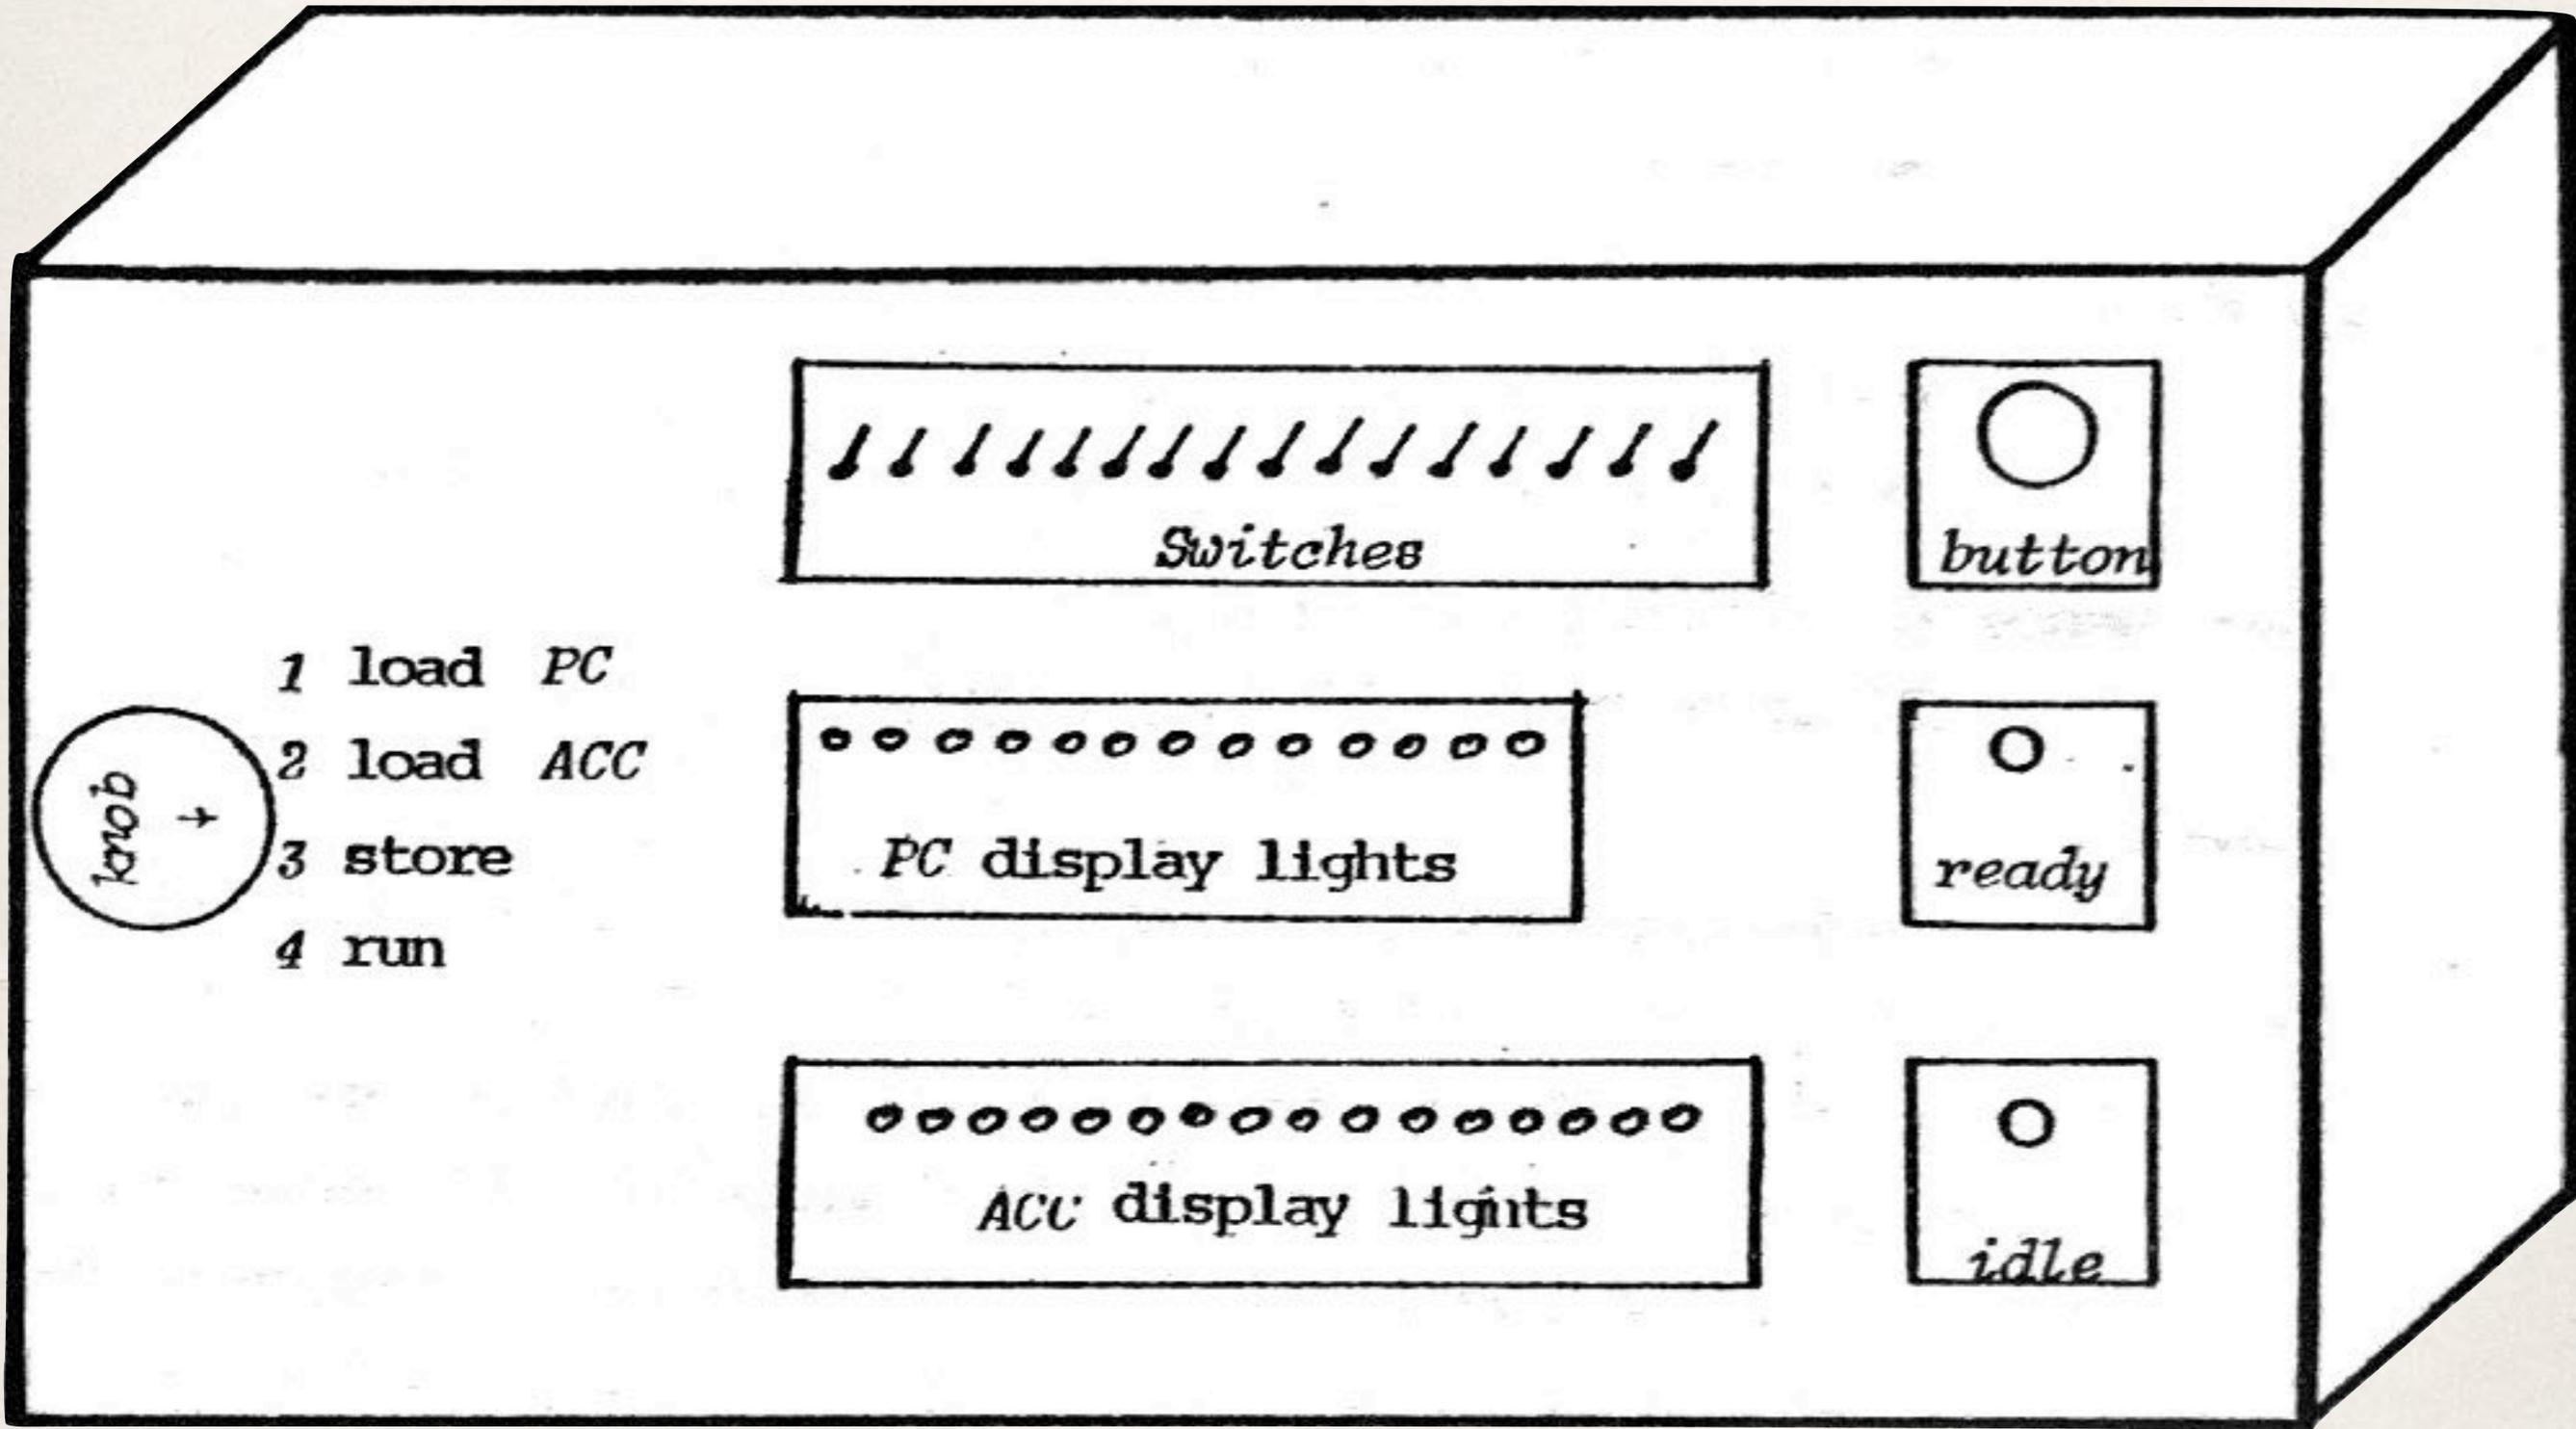
out((knob = k, button = b, switches = sw), (w, m, w₀, w₁, w₂, w₃, w₄, w₅))

= (pc = w₁, acc = w₂, ready = (w = 0 ∨ w = 5 + 1, 0), idle = (w = 0 + 1, 0))

next((knob = k, button = b, switches = sw), (w, m, w₀, w₁, w₂, w₃, w₄, w₅))

= (w = 0 + ((b + 1, 0) , m , w₀ , w₁ , w₂ , w₃ , w₄ , w₅) ,
 w = 1 + (0 , m , w₀ , w₁ , w₂ , w₃ , w₄ , w₅) ,
 w = 2 + (0 , m , w₀ , sw < 0 : 12 > , w₂ , w₃ , w₄ , w₅) ,
 w = 3 + (0 , m , w₀ , w₁ , sw , w₃ , w₄ , w₅) ,
 w = 4 + (7 , m , w₁ , w₁ , w₂ , w₃ , w₄ , w₅) ,
 w = 5 + (b + 0, 6) , m , w₀ , w₁ , w₂ , w₃ , w₄ , w₅) ,
 w = 6 + (8 , m , w₁ , w₁ , w₂ , w₃ , w₄ , w₅) ,
 w = 7 + (0 , m[w₂/w₁] , w₀ , w₁ , w₂ , w₃ , w₄ , w₅) ,
 w = 8 + (9 , m , w₀ , w₁ , w₂ , m(w₀) , w₄ , w₅) ,
 w = 9 + (w₃ < 13 : 15 > + 10 , m , w₀ , w₁ , w₂ , w₃ , w₄ , w₅) ,
 w = 10 + (0 , m , w₀ , w₁ , w₂ , w₃ , w₄ , w₅) ,
 w = 11 + (5 , m , w₀ , w₃ < 0 : 12 > , w₂ , w₃ , w₄ , w₅) ,
 w = 12 + ((w₂ = 0 + 11, 17) , m , w₀ , w₁ , w₂ , w₃ , w₄ , w₅) ,
 w = 13 + (19 , m , w₀ , w₁ , w₂ , w₃ , w₂ , w₅) ,
 w = 14 + (22 , m , w₀ , w₁ , w₂ , w₃ , w₂ , w₅) ,
 w = 15 + (24 , m , w₃ < 0 : 12 > , w₁ , w₂ , w₃ , w₄ , w₅) ,
 w = 16 + (25 , m , w₃ < 0 : 12 > , w₁ , w₂ , w₃ , w₄ , w₅) ,
 w = 17 + (18 , m , w₀ , w₁ , w₂ , w₃ , w₄ , w₁ < 0 : 12 > + 1) ,
 w = 18 + (6 , m , w₀ , w₅ < 0 : 12 > , w₂ , w₃ , w₄ , w₅) ,
 w = 19 + (20 , m , w₃ < 0 : 12 > , w₁ , w₂ , w₃ , w₄ , w₅) ,
 w = 20 + (21 , m , w₀ , w₁ , w₂ , w₃ , w₄ , w₄ + m(w₀)) ,
 w = 21 + (17 , m , w₀ , w₁ , w₅ , w₃ , w₄ , w₅) ,
 w = 22 + (23 , m , w₃ < 0 : 12 > , w₁ , w₂ , w₃ , w₄ , w₅) ,
 w = 23 + (21 , m , w₀ , w₁ , w₂ , w₃ , w₄ , w₄ - m(w₀)) ,
 w = 24 + (17 , m , w₀ , w₁ , m(w₀) , w₃ , w₄ , w₅) ,
 w = 25 + (17 , m[w₂/w₀] , w₀ , w₁ , w₂ , w₃ , w₄ , w₅))

even the industrial design!



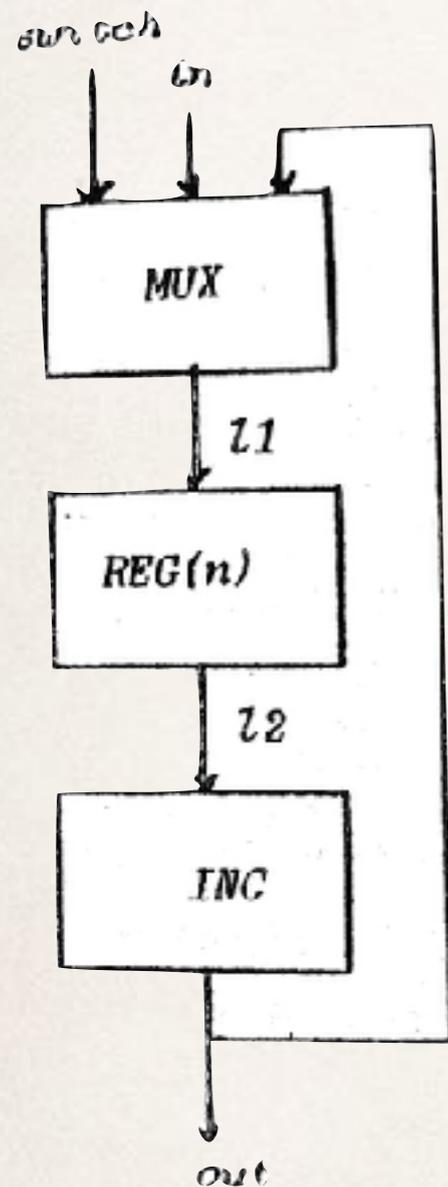
- ❖ Mike wrote *21 pages* on the computer alone
- ❖ 75 pages on a theory of hardware verification, from gates to computers
- ❖ Already many key insights by 1981, e.g. to treat *combinational* (stateless) devices like *sequential* ones.

Three ideas for modelling devices

- ❖ 1981: recursive domain equations (*too complicated*)
- ❖ 1983: Logic for Sequential Machines, or LSM: Mike's hardware formalism, loosely based on CCS (*too ad-hoc*)
- ❖ 1985: higher-order logic! And all devices as **relations**.

A concerted effort to *minimise* reliance on theory!

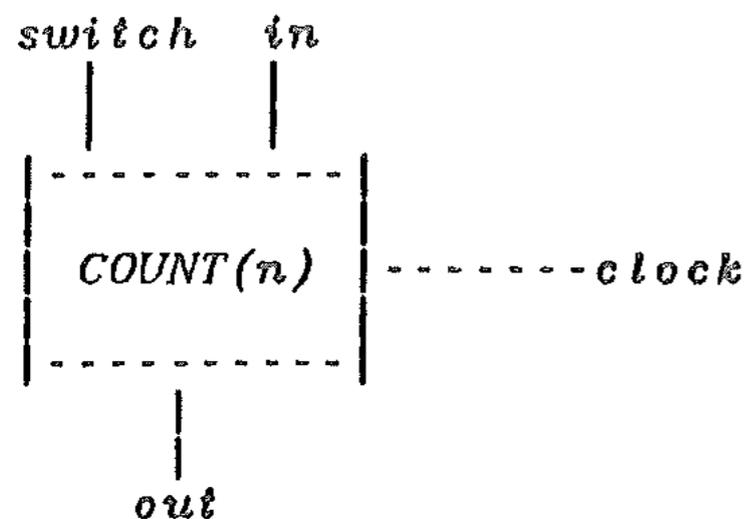
A counter using domains (1981)



$$COUNT(n) = \lambda(\text{switch}, in). (\text{out} = n + 1), COUNT(\text{switch} + in, n + 1)$$

Fig. II. Behaviour of $COUNT(n)$

A counter in LSM (1983)



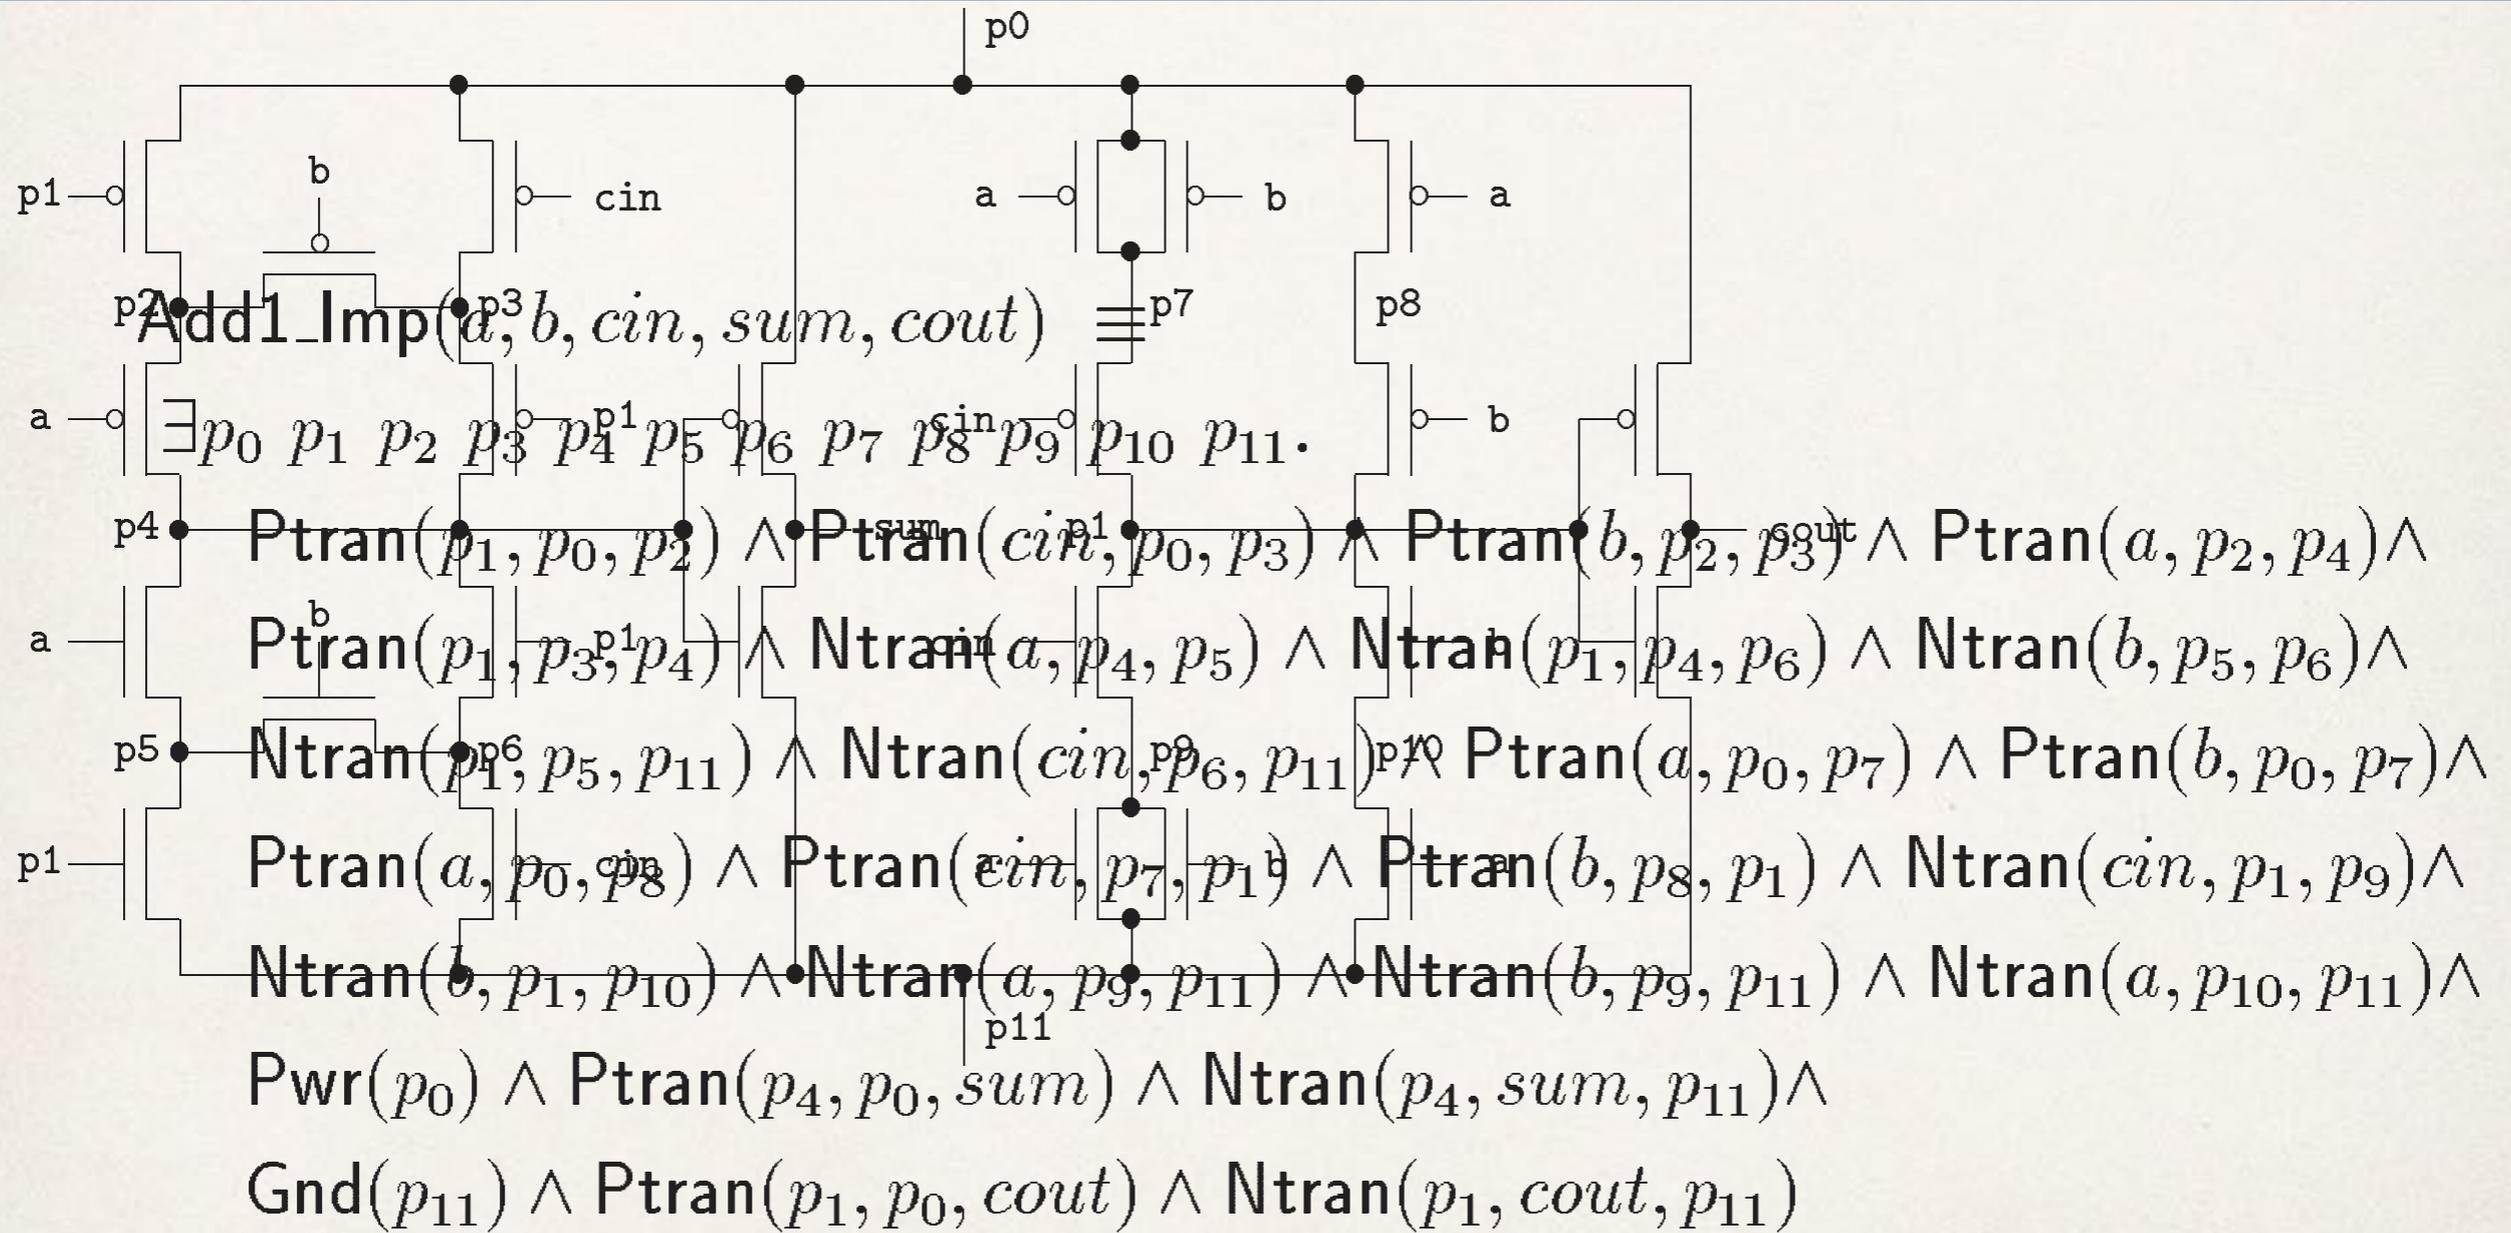
When the counter is clocked, the new value of the state variable n becomes $n+1$ (i.e. the old value plus one) if *false* is being input on line *switch*, otherwise it becomes the value input on line *in*. We can express this by:

$$CLOCK(n) \rightarrow CLOCK(\text{switch} \rightarrow \text{in} \mid n+1)$$

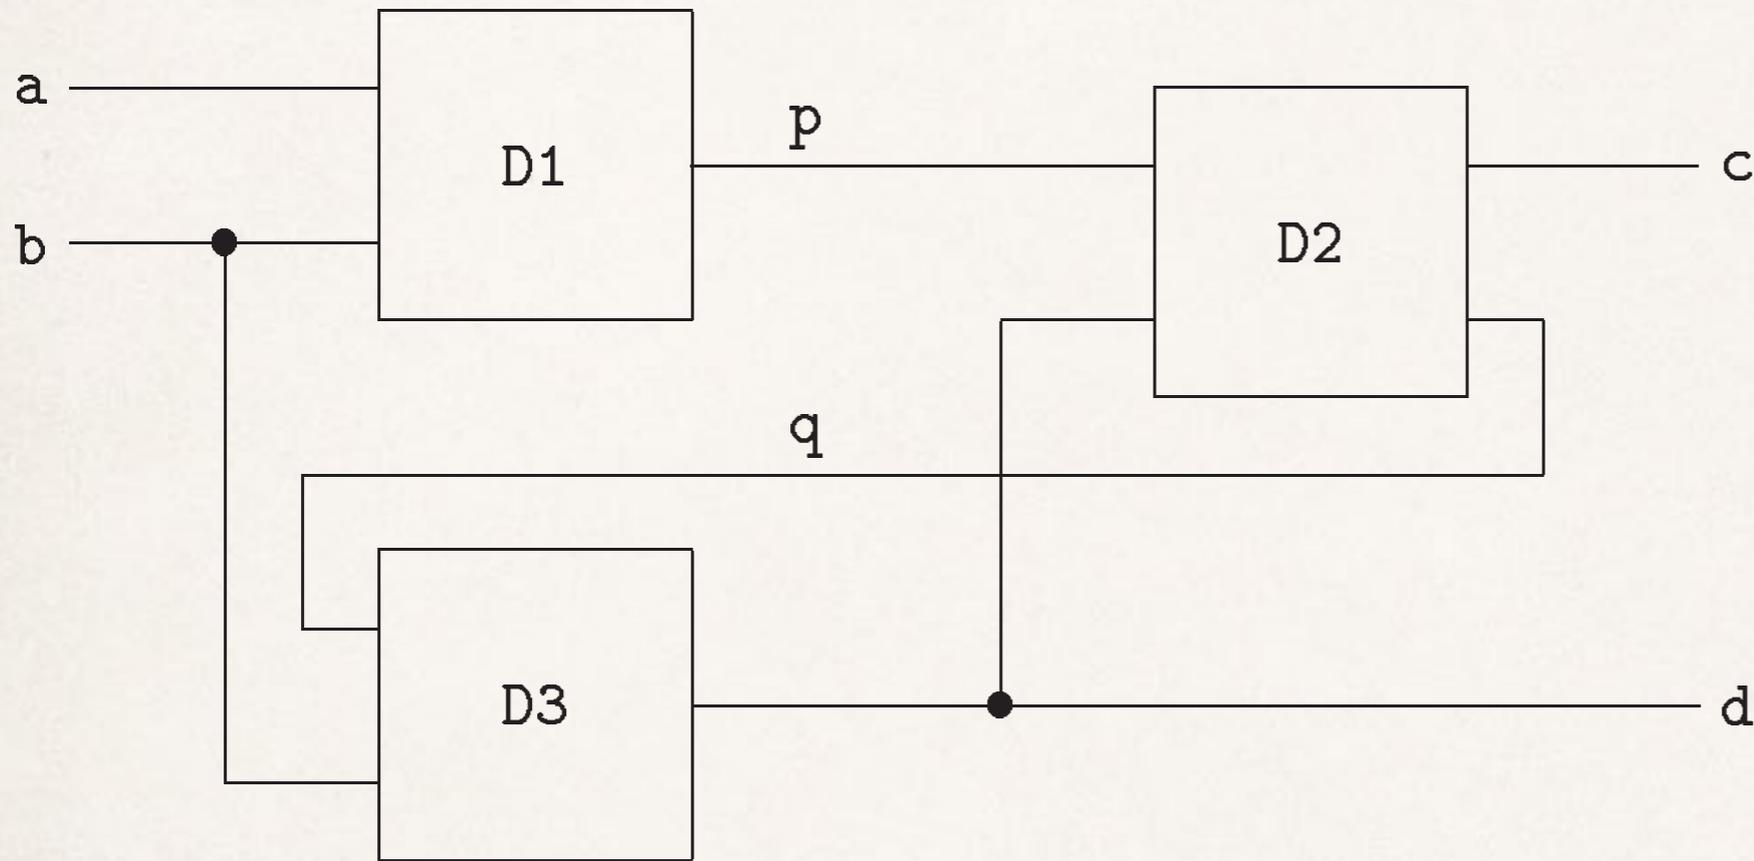
In LSM the behaviour of the counter would be specified by the formula

$$COUNT(n) == dev\{\text{switch}, \text{in}, \text{out}\}. \{\text{out}=n\}; COUNT(\text{switch} \rightarrow \text{in} \mid n+1)$$

A CMOS full adder in HOL (1985)



The insight that devices are *relations*



$$D(a, b, c, d) \equiv \exists p q. D_1(a, b, p) \wedge D_2(p, d, c) \wedge D_3(q, b, d)$$

HO Logic was a radical choice!

“Unlike first-order logic and some of its less baroque extensions, second and higher-order logic have no coherent well-established theory; the existent material consisting merely of scattered remarks quite diverse with respect to character and origin.”

(Van Benthem and Kees Doets, 1983.)

And we need a type of n -bit words, so we need dependent types, right?

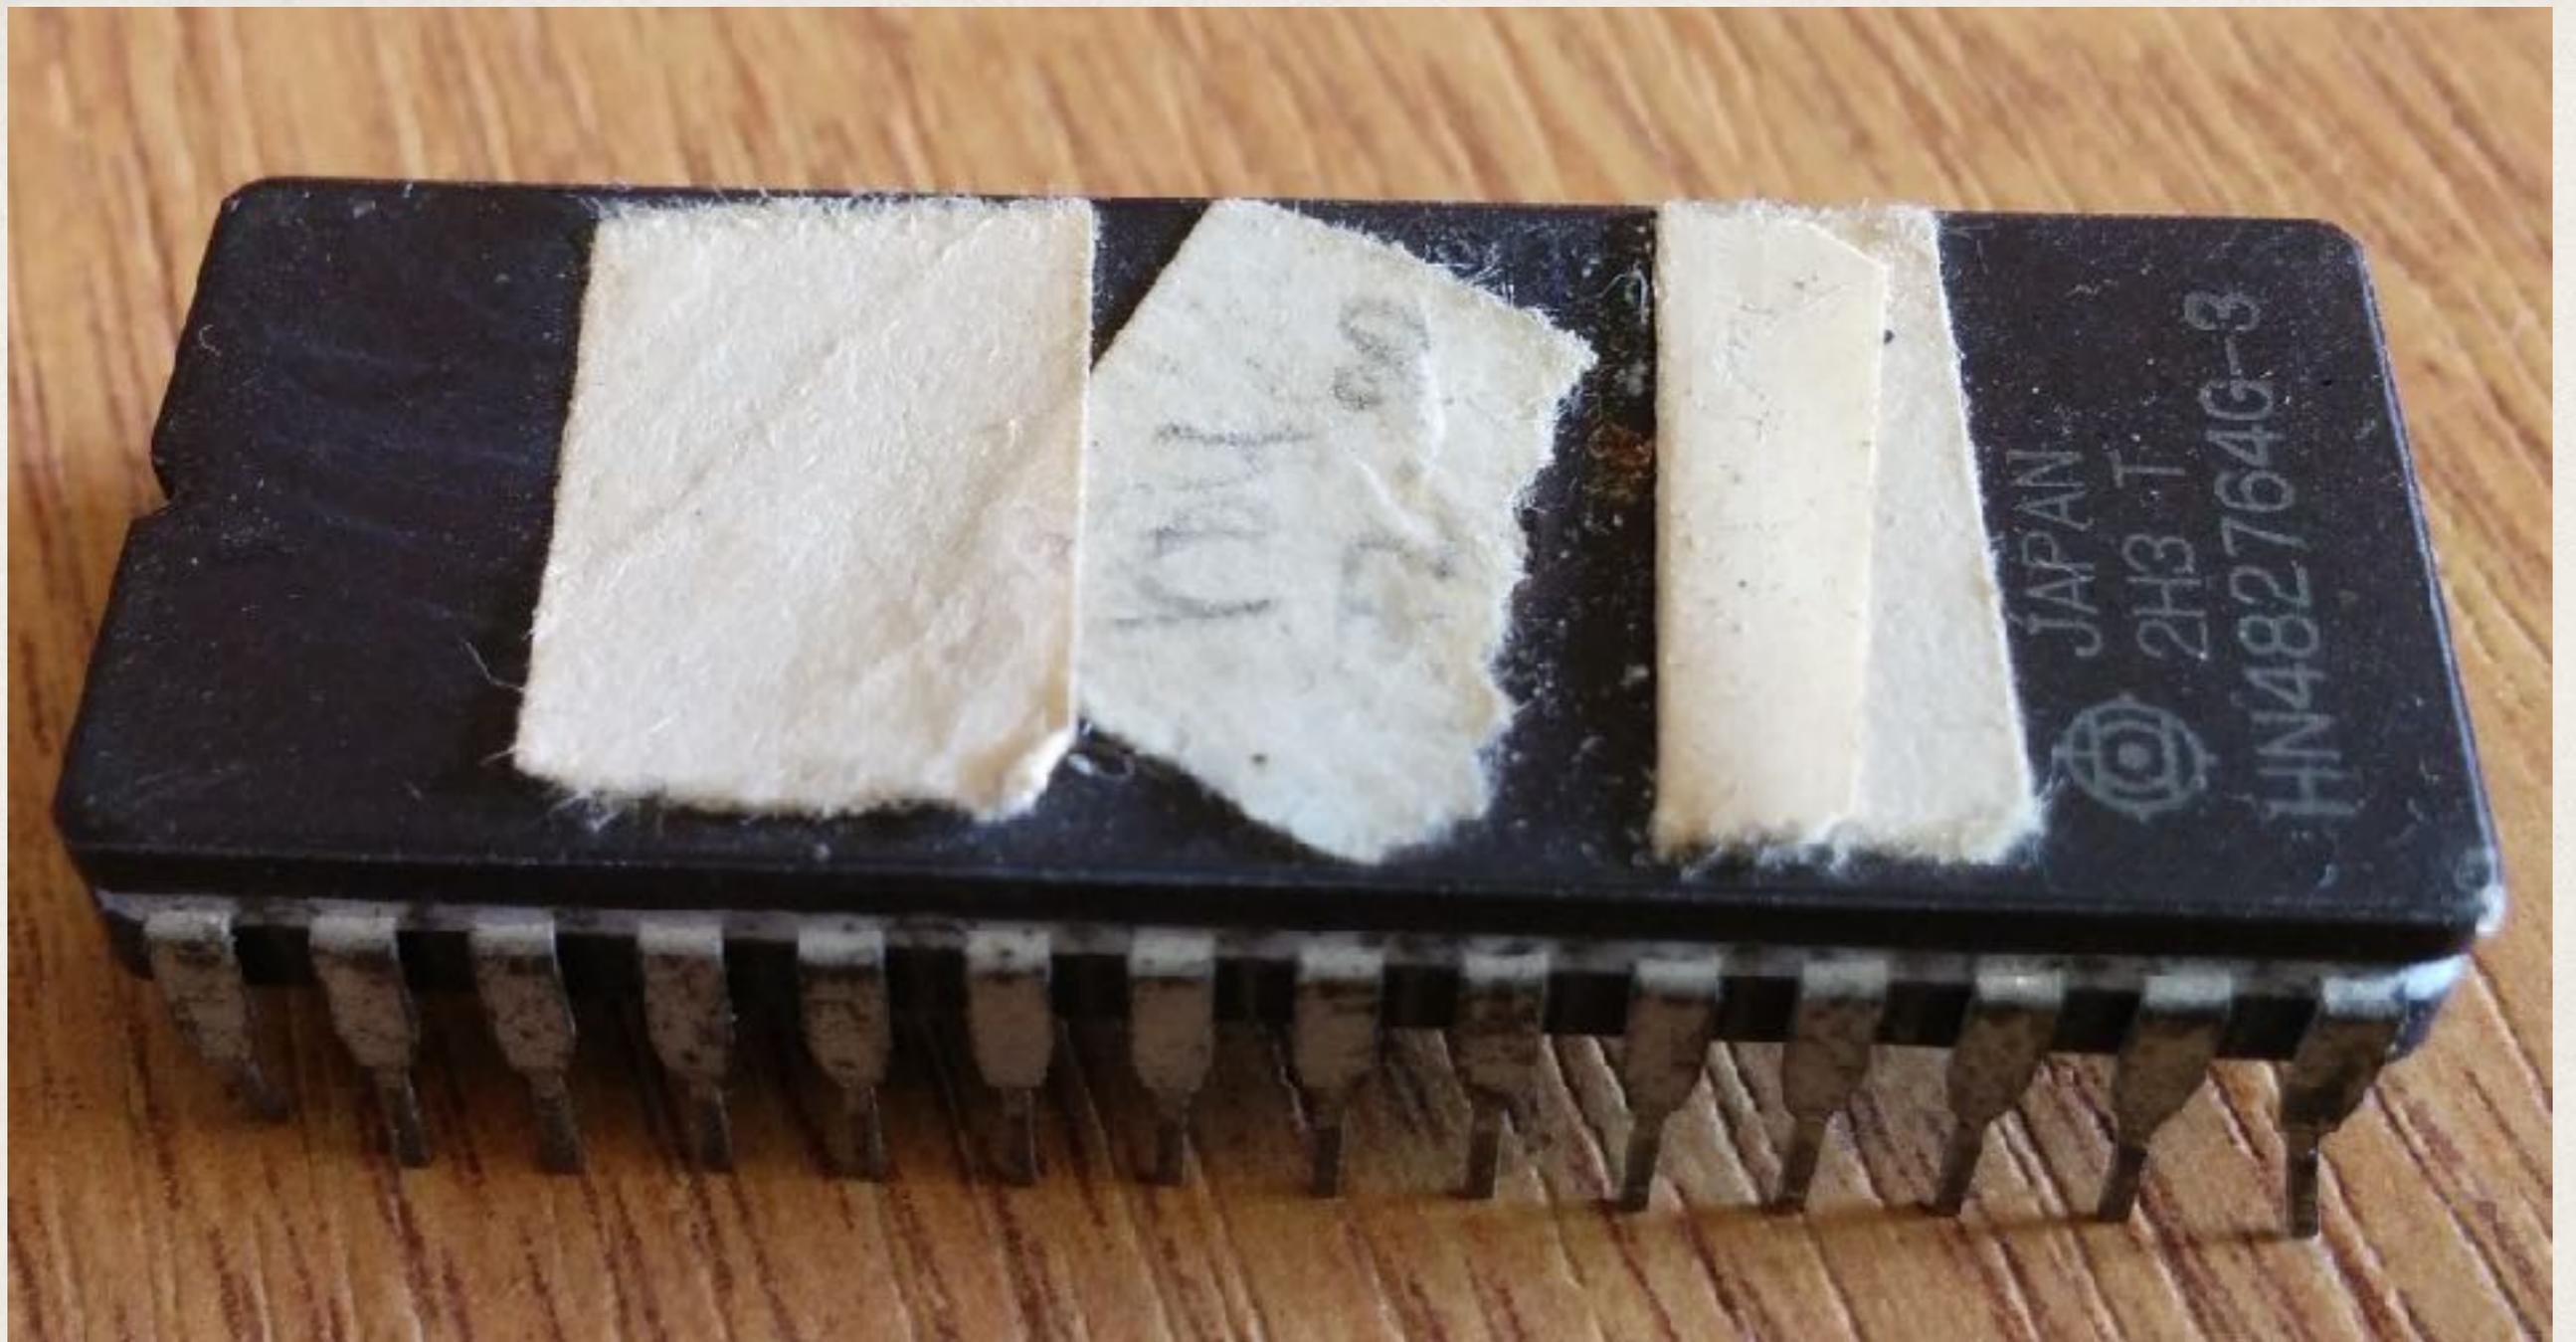
WRONG.

Verifying Mike's computer

- ❖ In 1983, using LCF_LSM

“The entire specification and verification described here took several months, but this includes some extending and debugging of LCF_LSM ... it would take me two to four weeks to do another similar exercise now. The complete proof requires several hours CPU time on a 2 megabyte Vax 750.”

- ❖ In 1986, with Jeff Joyce, Graham Birtwistle, using HOL
- ❖ And — under the name *Tamarack* — by many others!



It was even fabricated!

Verifying a *real* computer: VIPER

- ❖ designed by a UK defence lab for military purposes
- ❖ specified by a series of abstract layers
 - ❖ equivalence of the *top three* proved by Avra Cohn
 - ❖ controversy due to exaggerated claims made by the chip's marketers

Hardware verification went from plan
(1981) to realisation (1989) in eight years!

Some papers from that era

Mike Gordon. Proving a computer correct with the LCF LSM hardware verification system (1983).

Avra Cohn. A proof of correctness of the VIPER microprocessor: The first level (1987).

Avra Cohn. Correctness properties of the Viper block model: The second level (1988).

Brian T. Graham. *The SECD Microprocessor, A Verification Case Study*. Kluwer, 1992.

“Verification involves a pair of models that bear an uncheckable and possibly imperfect relation to the *intended design* and to the *actual device*.”

–*Avra Cohn, 1989*

So what next?

“a long term project on verifying combined hardware/software systems by mechanized formal proof”

Mike Gordon, “Mechanizing programming logics in higher order logic”, 1989

- ❖ Mike *derived* Hoare logic in HOL from the operational semantics of a programming language
- ❖ supporting the illusion by pretty-printing
- ❖ the first **shallow embedding** of one formalism within another

Whooshing to the year 2000...

Lots of PhDs on floating point arithmetic,
process calculi, BDDs, etc., etc....

Verifying the ARM6 processor

- ❖ a “commercial off-the-shelf” design
 - ❖ joint project with Graham Birtwistle at Leeds
 - ❖ verification by Anthony Fox at Cambridge
- ❖ *the ARM6 microarchitecture implements its ISA*
- ❖ *And we have a complete formal spec of this machine*

Verification and assembly language

(Magnus Myreen, working with Mike)

- ❖ Hoare-style logics for assembly languages
- ❖ decompilation of assembler to HOL (for 3 machines!)
- ❖ proof-producing translation from HOL to assembler

“verifying combined hardware/software systems”

Culminating in CakeML

A dialect of Standard ML, with semantics formalised in HOL

The compiler backend can generate code for 5 different architectures

Source code can be written directly or translated from HOL

Bootstrapped via compilation within HOL4, yielding a
“verified binary that provably implements the compiler itself”

The work of Ramana Kumar, Magnus Myreen, Scott Owens, etc.

How Mike accomplished so much

- ❖ he ignored “hot topics” to pursue an original plan
- ❖ ... and talked to “the enemy” across the Department
- ❖ ... to learn another subject *really well*
- ❖ while using his own knowledge of theory (denotational semantics and CCS)

Learn your application *thoroughly*

E.g. *cryptographic protocols*: such a **simple** field

... so much flawed research

... leading to more bad research

Because people weren't
learning from *security* experts!

Rely on robust tools and theory

- ❖ LCF provided a good verification engine
- ❖ Higher-order logic was old, solid theory (1940!)
- ❖ [Standard ML was emerging; Mike didn't use it!]

Type theory would take many years to settle down and would have been a distraction back in 1985.

Where are we today?

- ❖ LCF architectures dominate the landscape
- ❖ ... using higher-order logic or its extensions
- ❖ hardware verification is done extensively in industry
- ❖ academic research continues to push forward

But modern processors are *still*
too complex to verify in full!

Mike Gordon

1948–2017

