Perspectives from Four Decades of Chip Design

Paul Scheidt

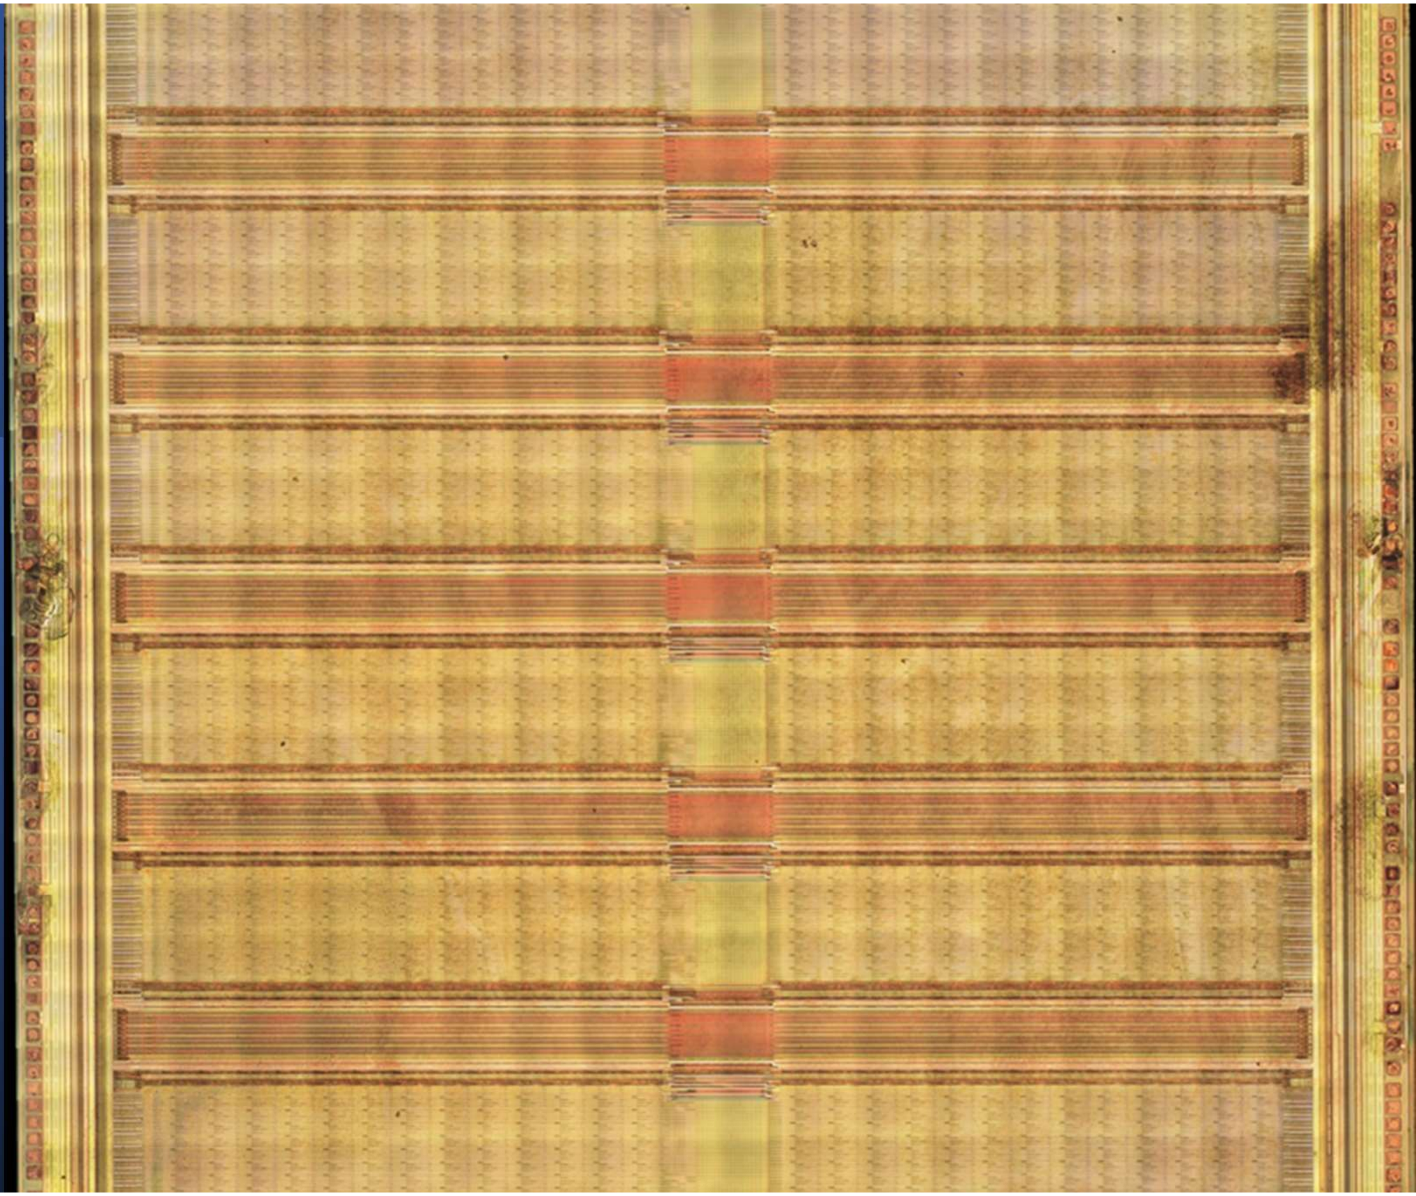

Thanks to my corporate sponsor!

Geographical Constraints

Communications across a sparsely populated land

Cross-Country Connectivity

Trans-Canada Microwave Relay Network

7000 km
6 Time Zones

https://en.wikipedia.org/wiki/Trans_Canada_Microwave

Continental Interconnect

United States Long Line Network

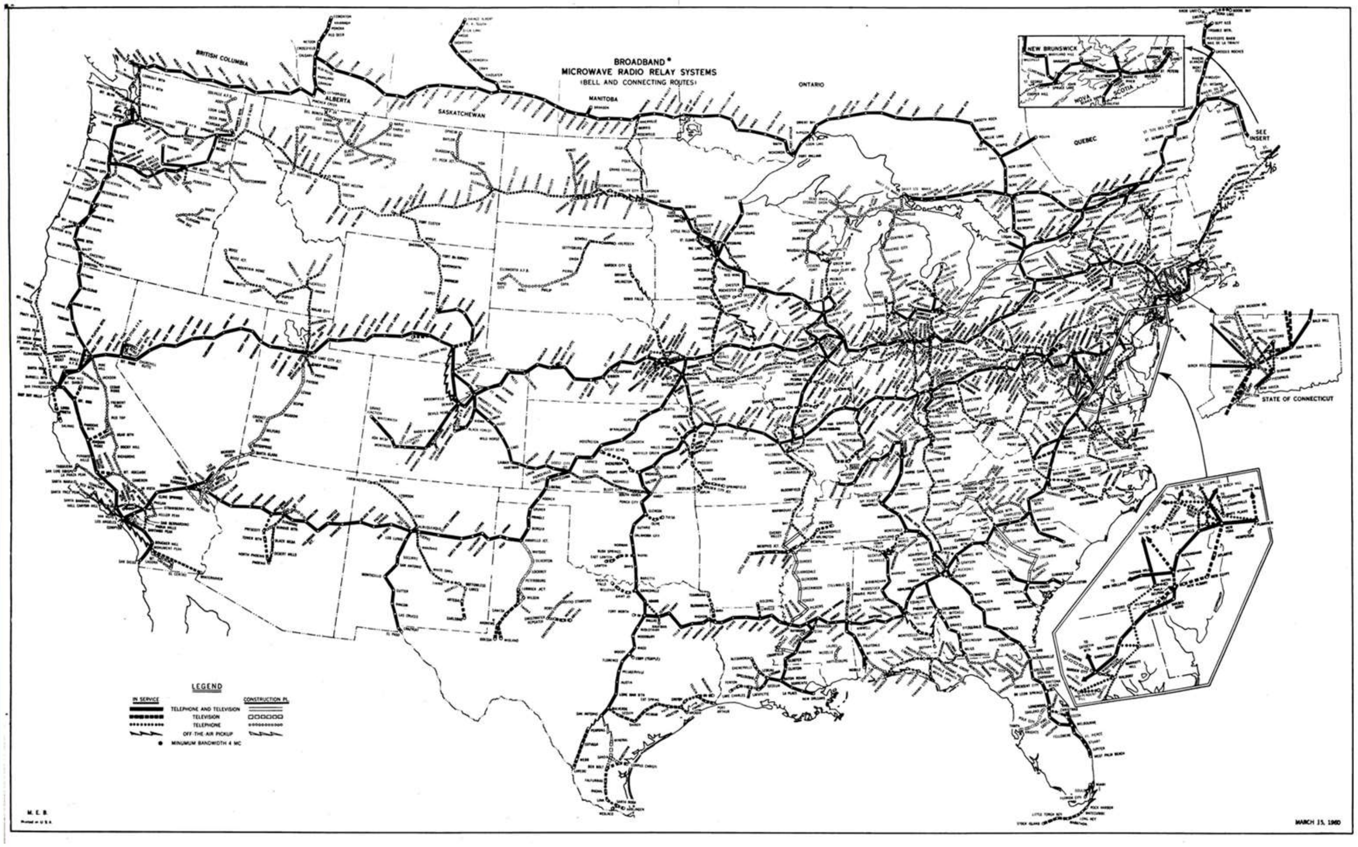

https://www.long-lines.net/places-routes/maps/MW6003.jpg

The Beginning

Point of Sale Terminal

Early Electronic Commerce

https://www.istockphoto.com/photo/old-german-telephone-booth-gm458732917-20202153

Smart Card Concept Prototype

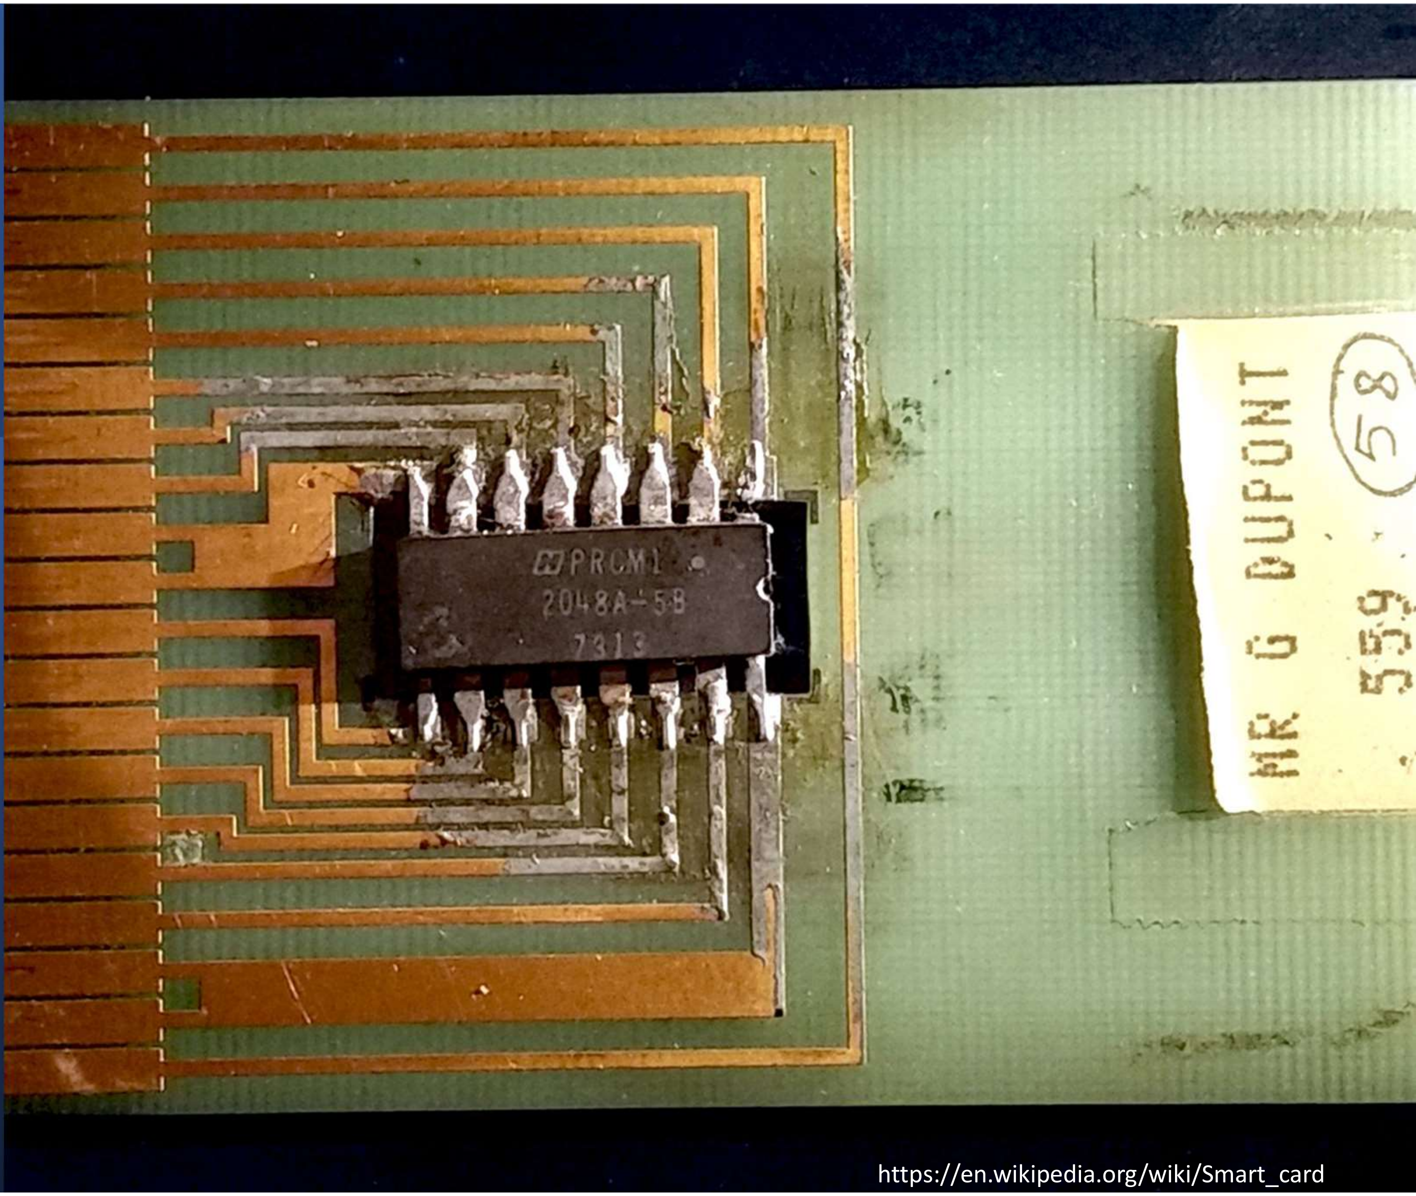

https://en.wikipedia.org/wiki/Smart_card

Old School Chip Emulator

Racks full of Wire Wrapped Cards

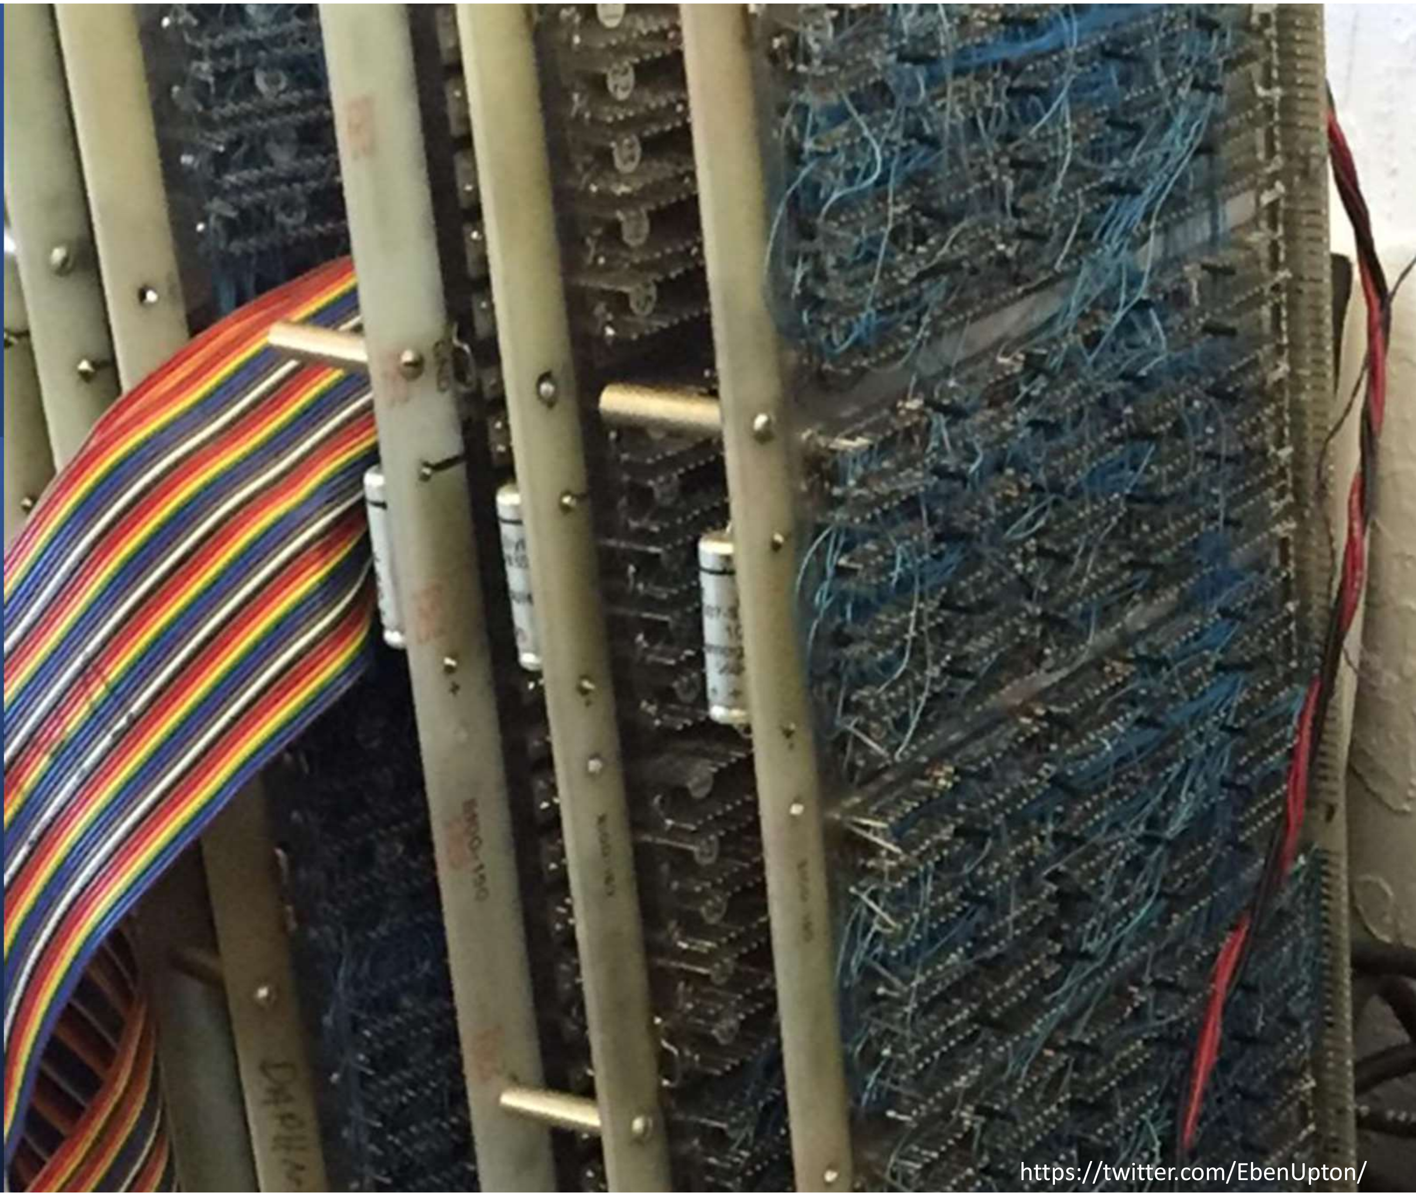

https://twitter.com/EbenUpton/

HAPS Prototype

Zebu Emulator

Leverage FPGAs

Hand Layout

Hand drawn on Mylar and digitized

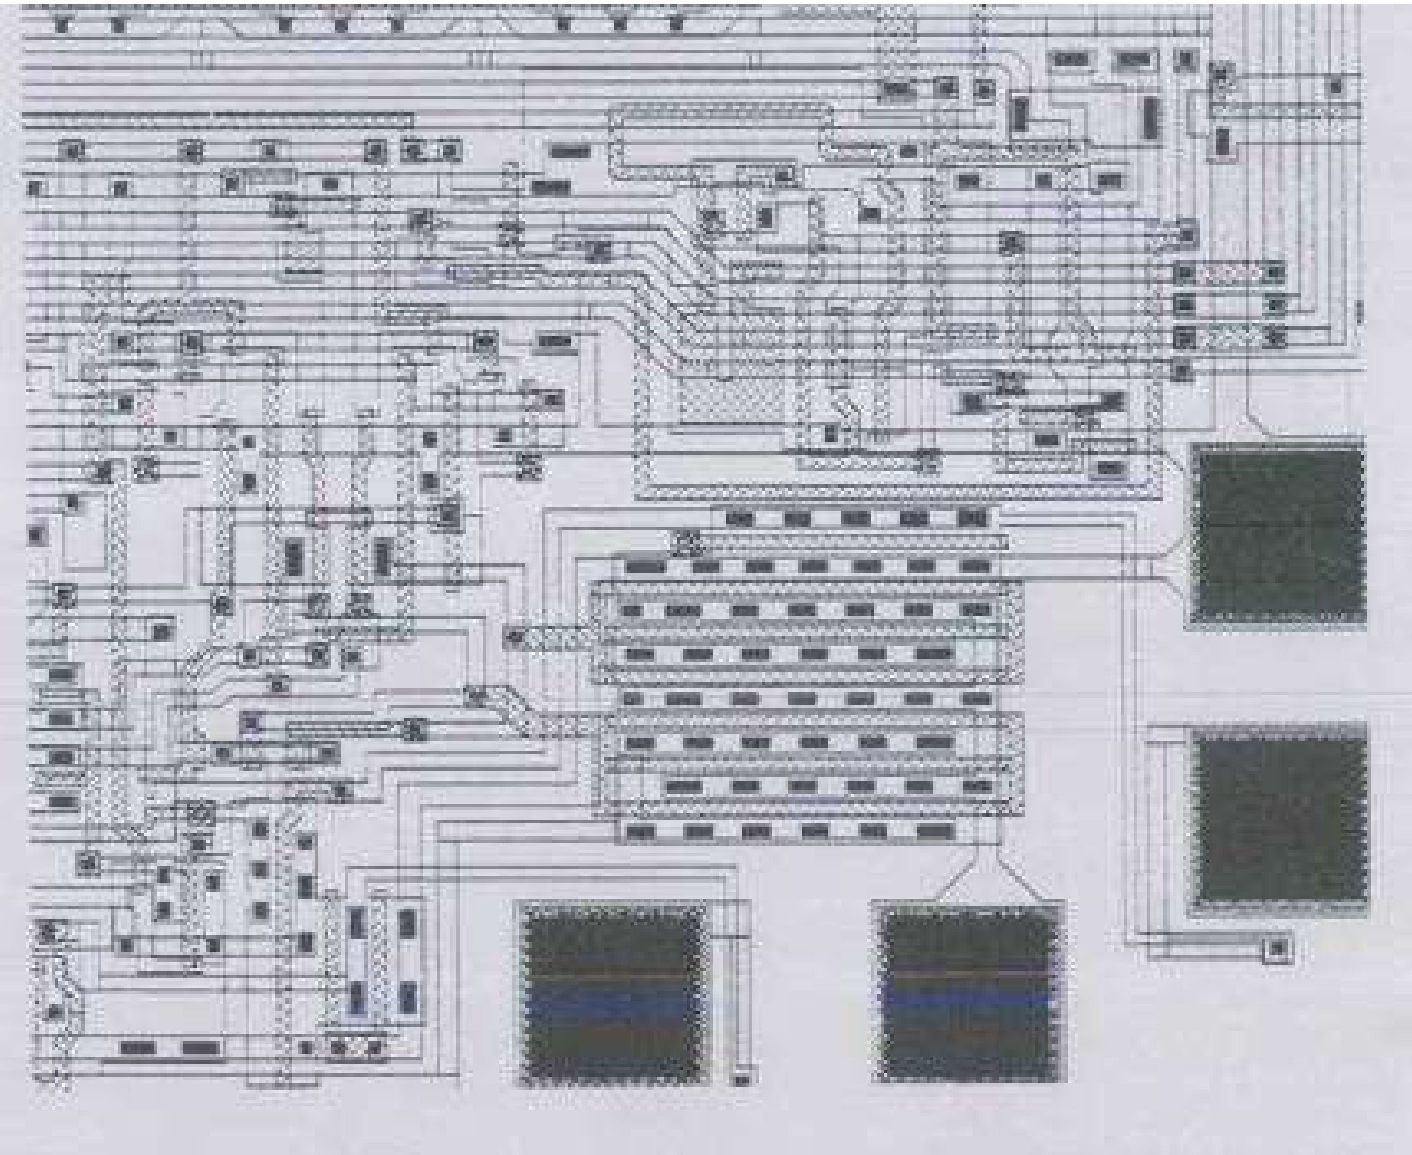

Courtesy Calma, Inc.

VLSI CAD Workstation 1980s

Layout Digitization

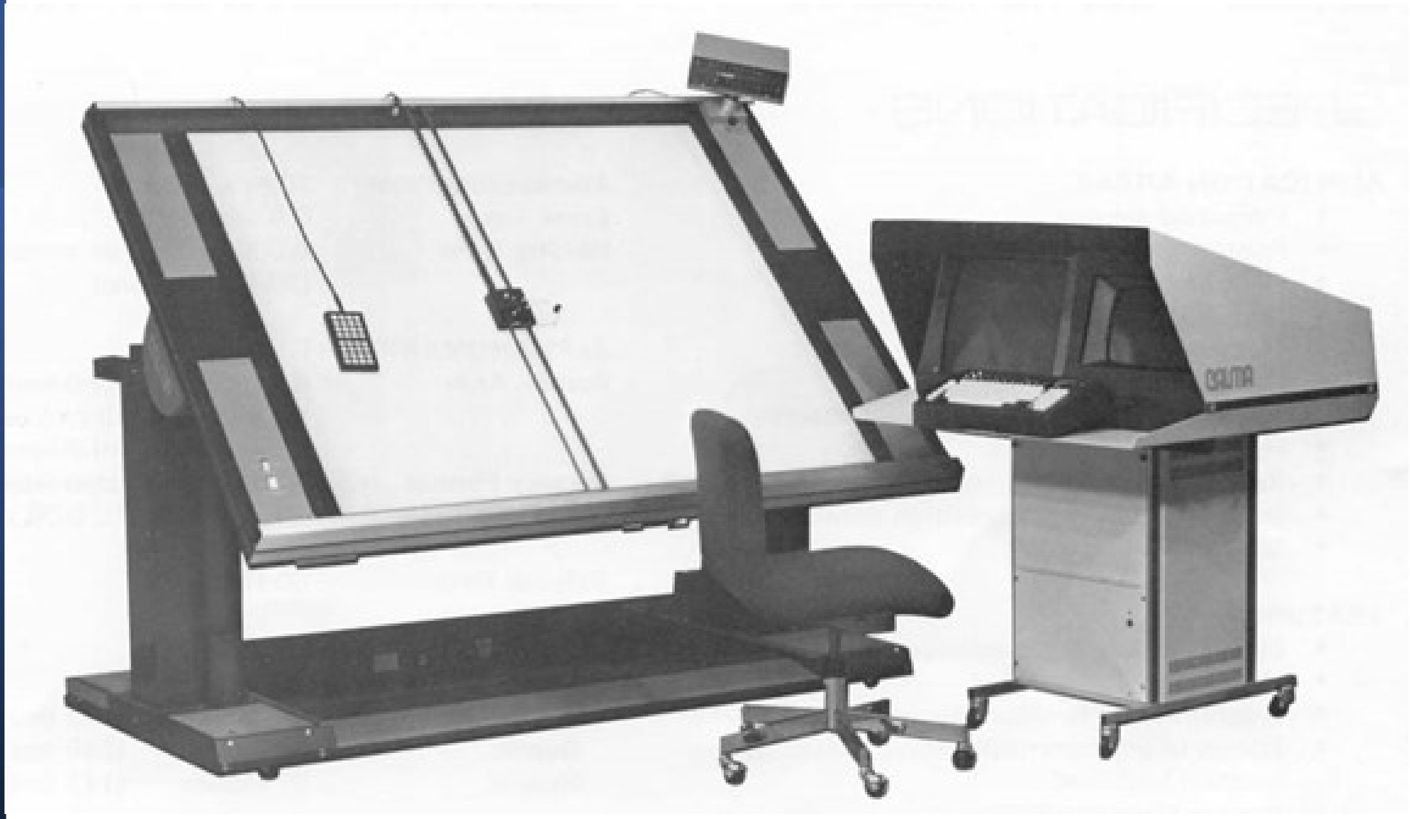

https://www.shapr3d.com/history-of-cad/calma

Siemens Secure EEPROM

2.5 micron Depletion Load NMOS

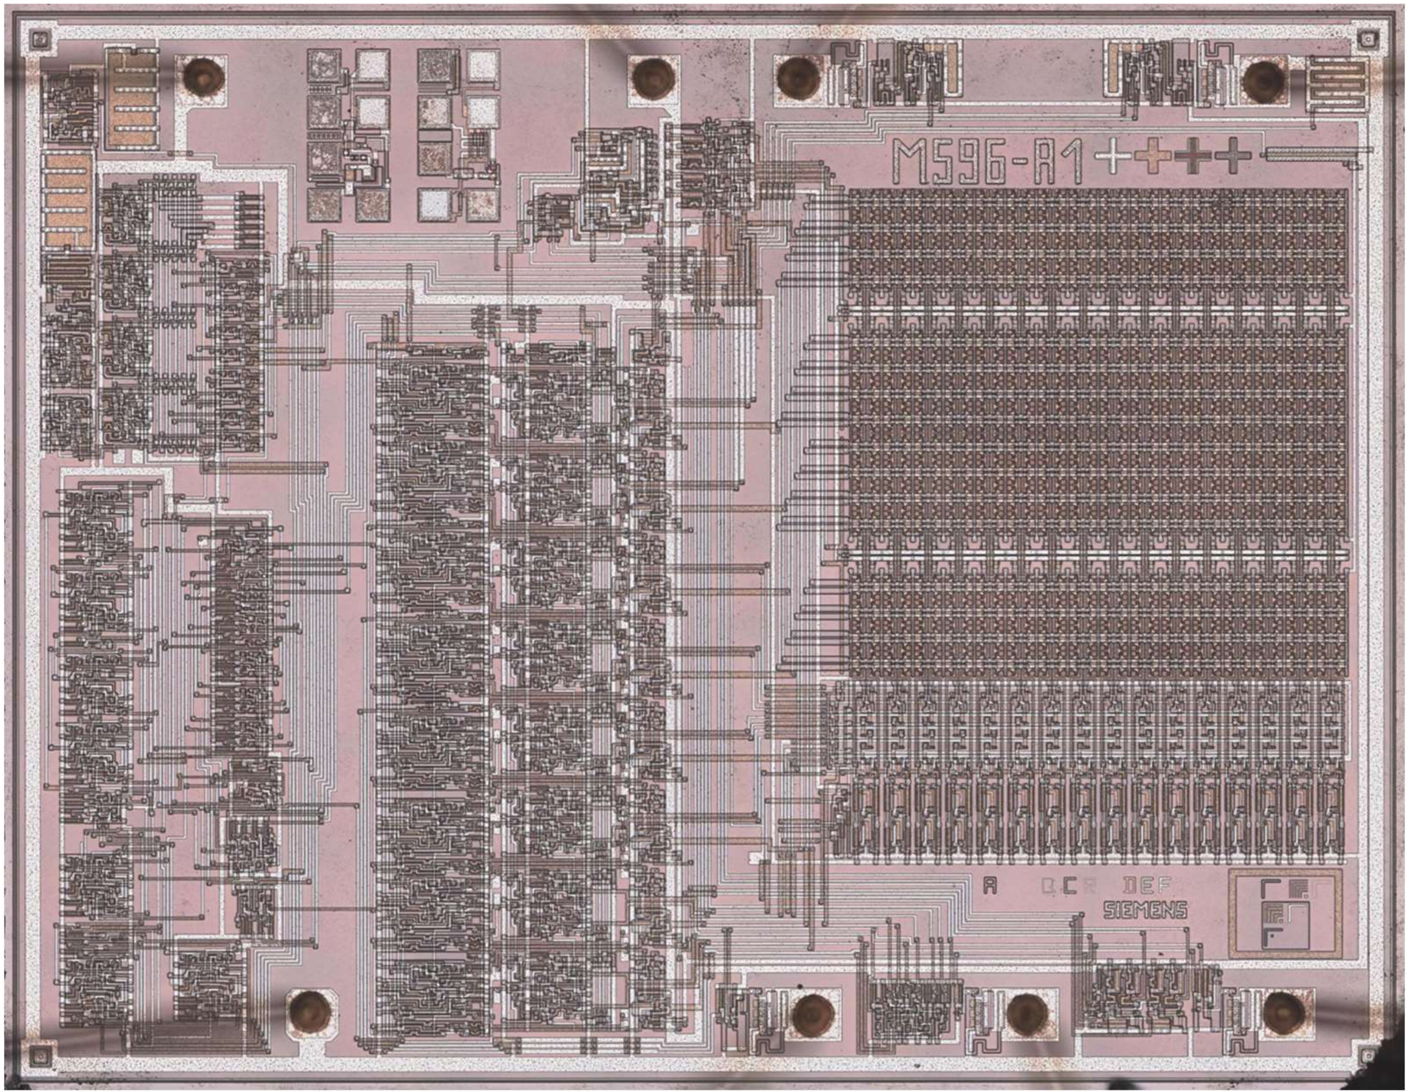

Source: Paul Scheidt – scanned by John McMaster
https://siliconpr0n.org/map/siemens/m596-a1/mcmaster_mz_mit20x/

MOSFET Technology

early-mid 1980s

NMOS Depletion Load Logic

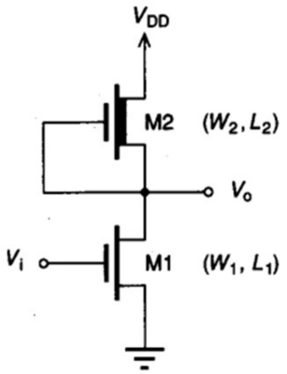

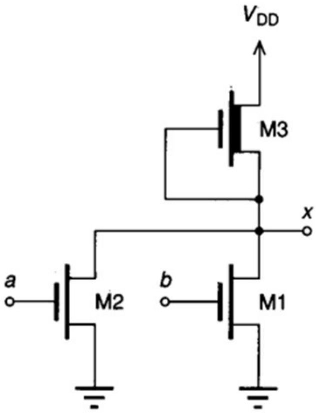

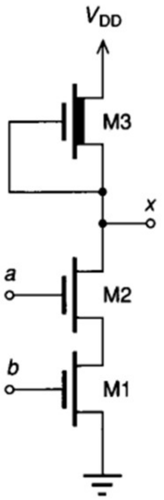

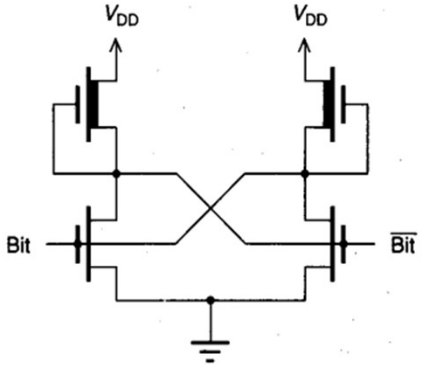

Attempts to Go Faster & Smaller

Dynamic
Domino Logic

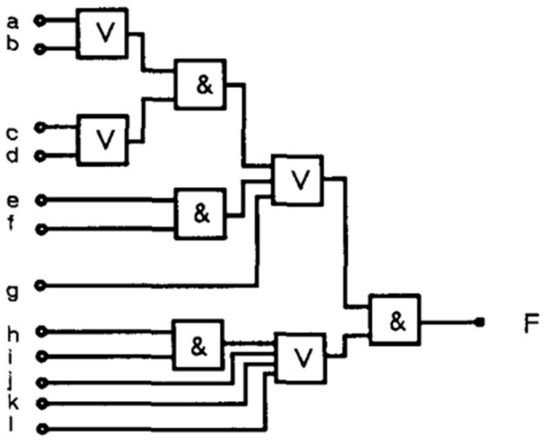

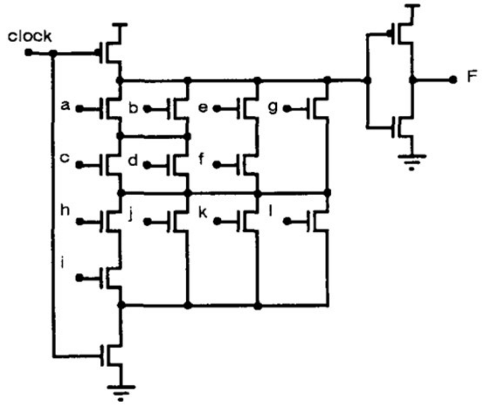

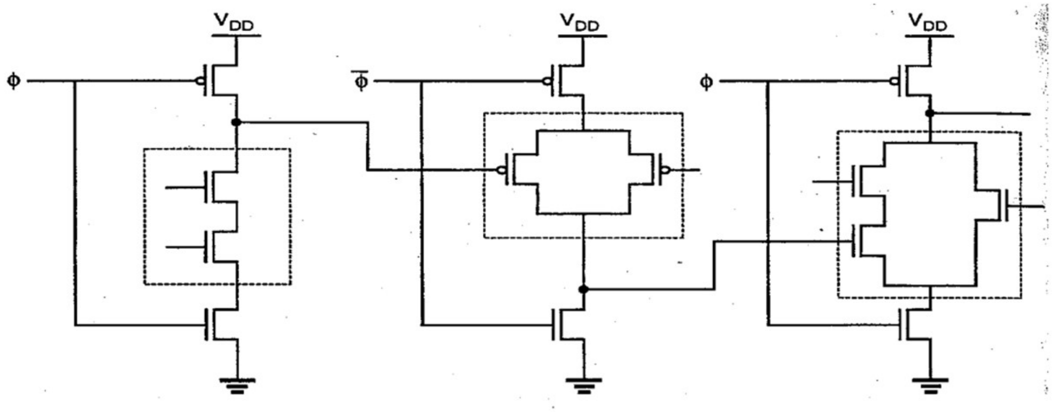

IEEE Journal of Solid State Circuits Vol 25, No. 2, April 1990

Dynamic Logic soon relegated to Niche

Static CMOS won the race!

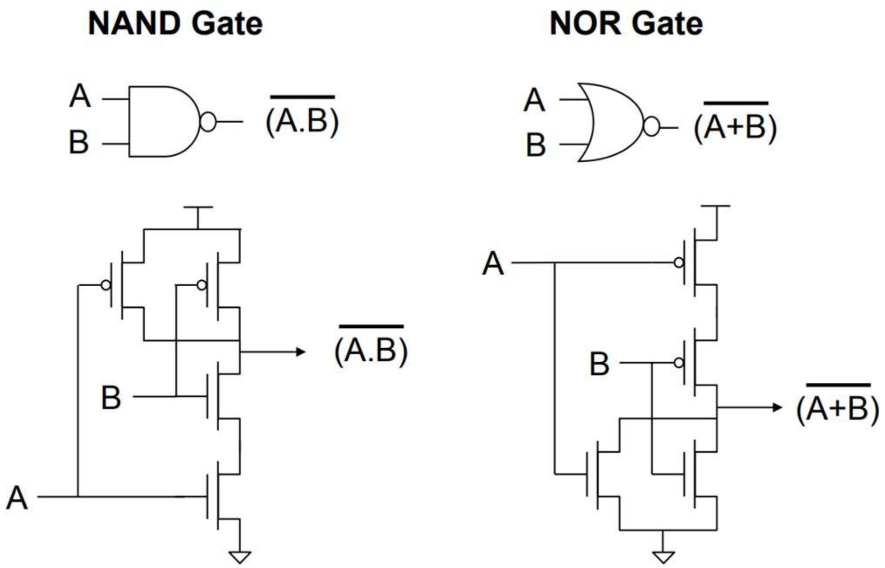

Yet…
Dynamic Logic is Relevant Again!

ST-TDPL
Self-Timed Three-Phase Dual-Rail Precharge Logic

Side Channel Resistant

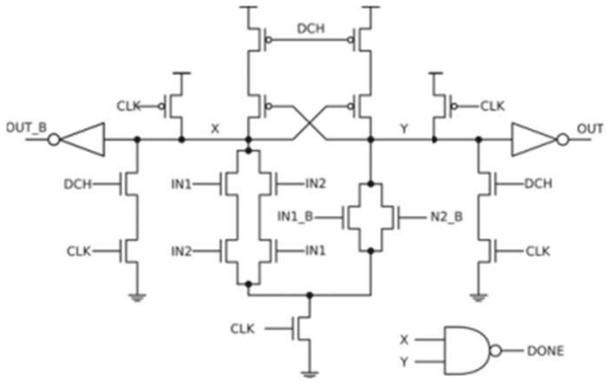

(a) NAND

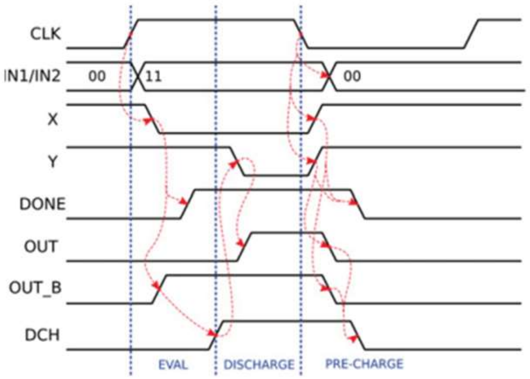

Fig. 7: Timing diagram of a ST-TDPL NAND gate for an input transition of 00 to 11.

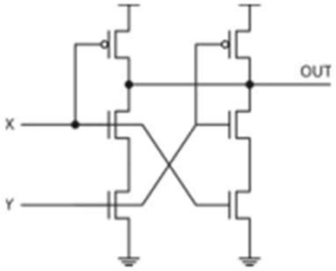

Fig. 6: DONE signal generating balanced NAND gate.

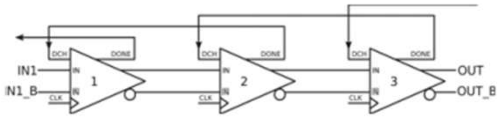

Nail Etkin, "A DPA-Resistant Self-Timed Three-Phase Dual-Rail Pre-Charge Logic Family", 2015 IEEE HOST

MOS Current Mode Logic (MCML)

Reminiscent of Emitter Coupled Logic (ECL)

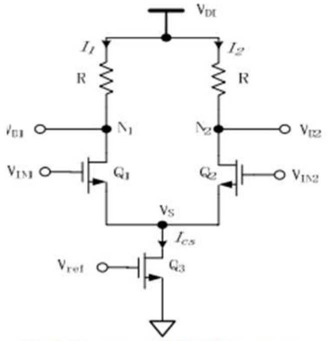

Fig 1. Structure of MCML inverter

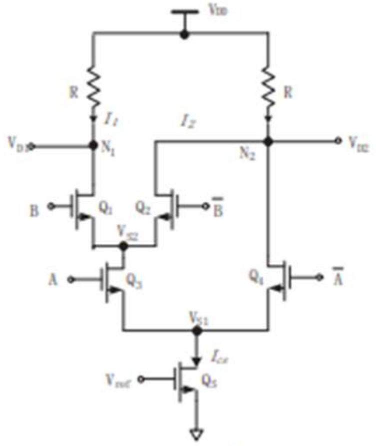

(a) AND2

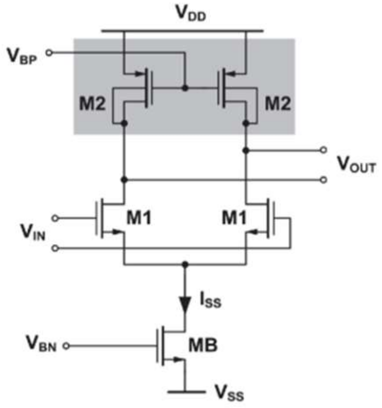

Fig 1 CML Inverter Circuit

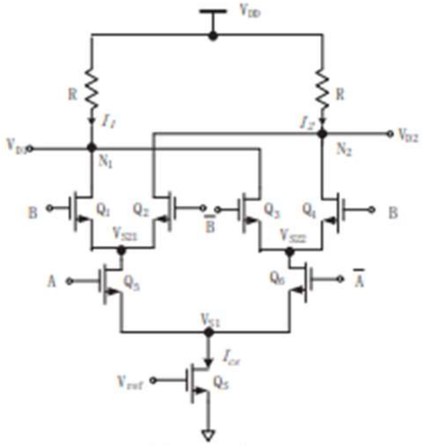

(c) XOR2

https://ieeexplore.ieee.org/stamp/stamp.jsp?tp=&arnumber=8529461

https://ieeexplore.ieee.org/stamp/stamp.jsp?tp=&arnumber=6230289

Advancing Telecom Systems

Northern Telecom DMS-100

Texas Instruments TMS32010

Start of the DSP Revolution

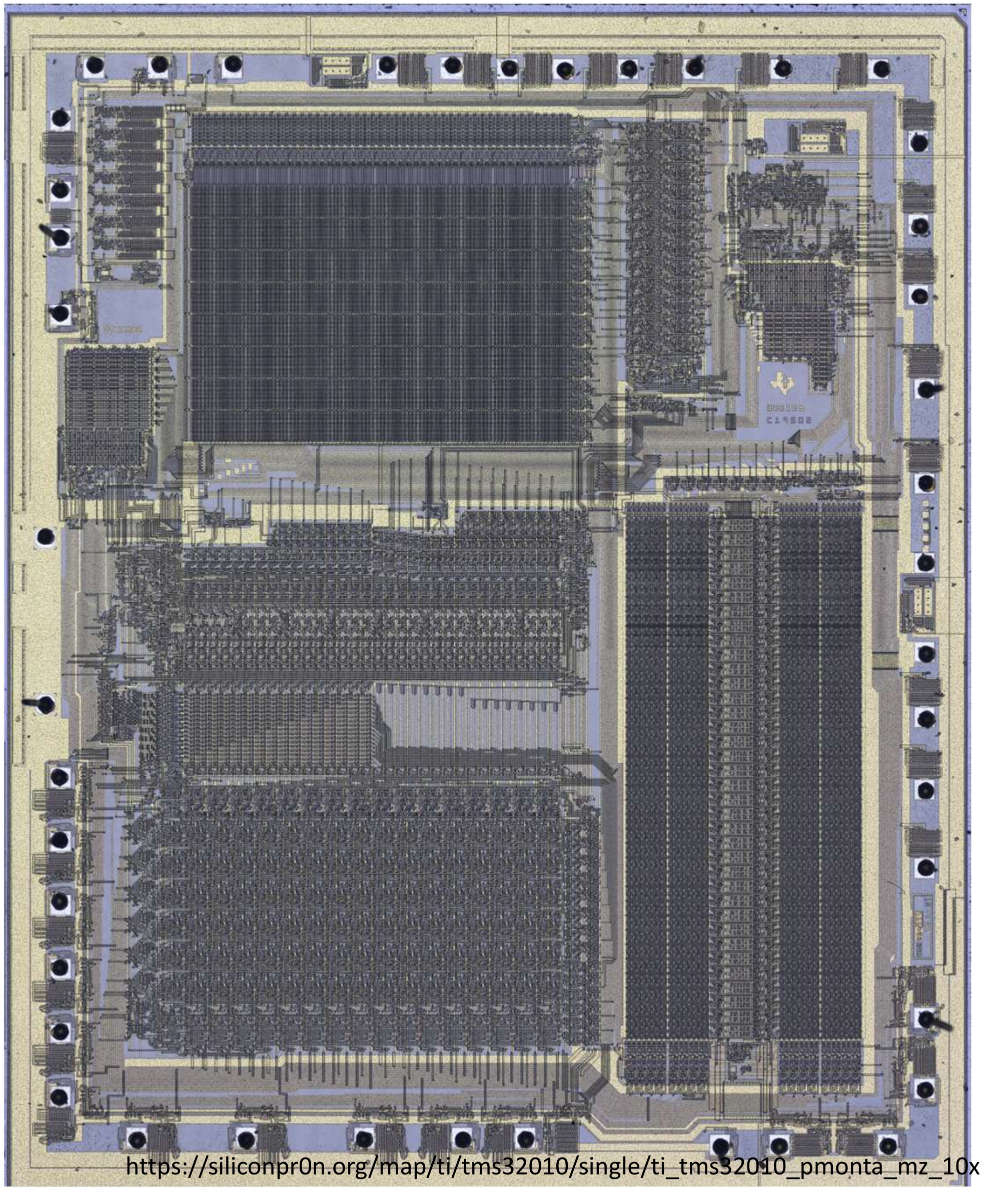

https://siliconpr0n.org/map/ti/tms32010/single/ti_tms32010_pmonta_mz_10x.jpg

Mixed-Signal Explorations

Minimizing Analog with Σ-Δ Modulation

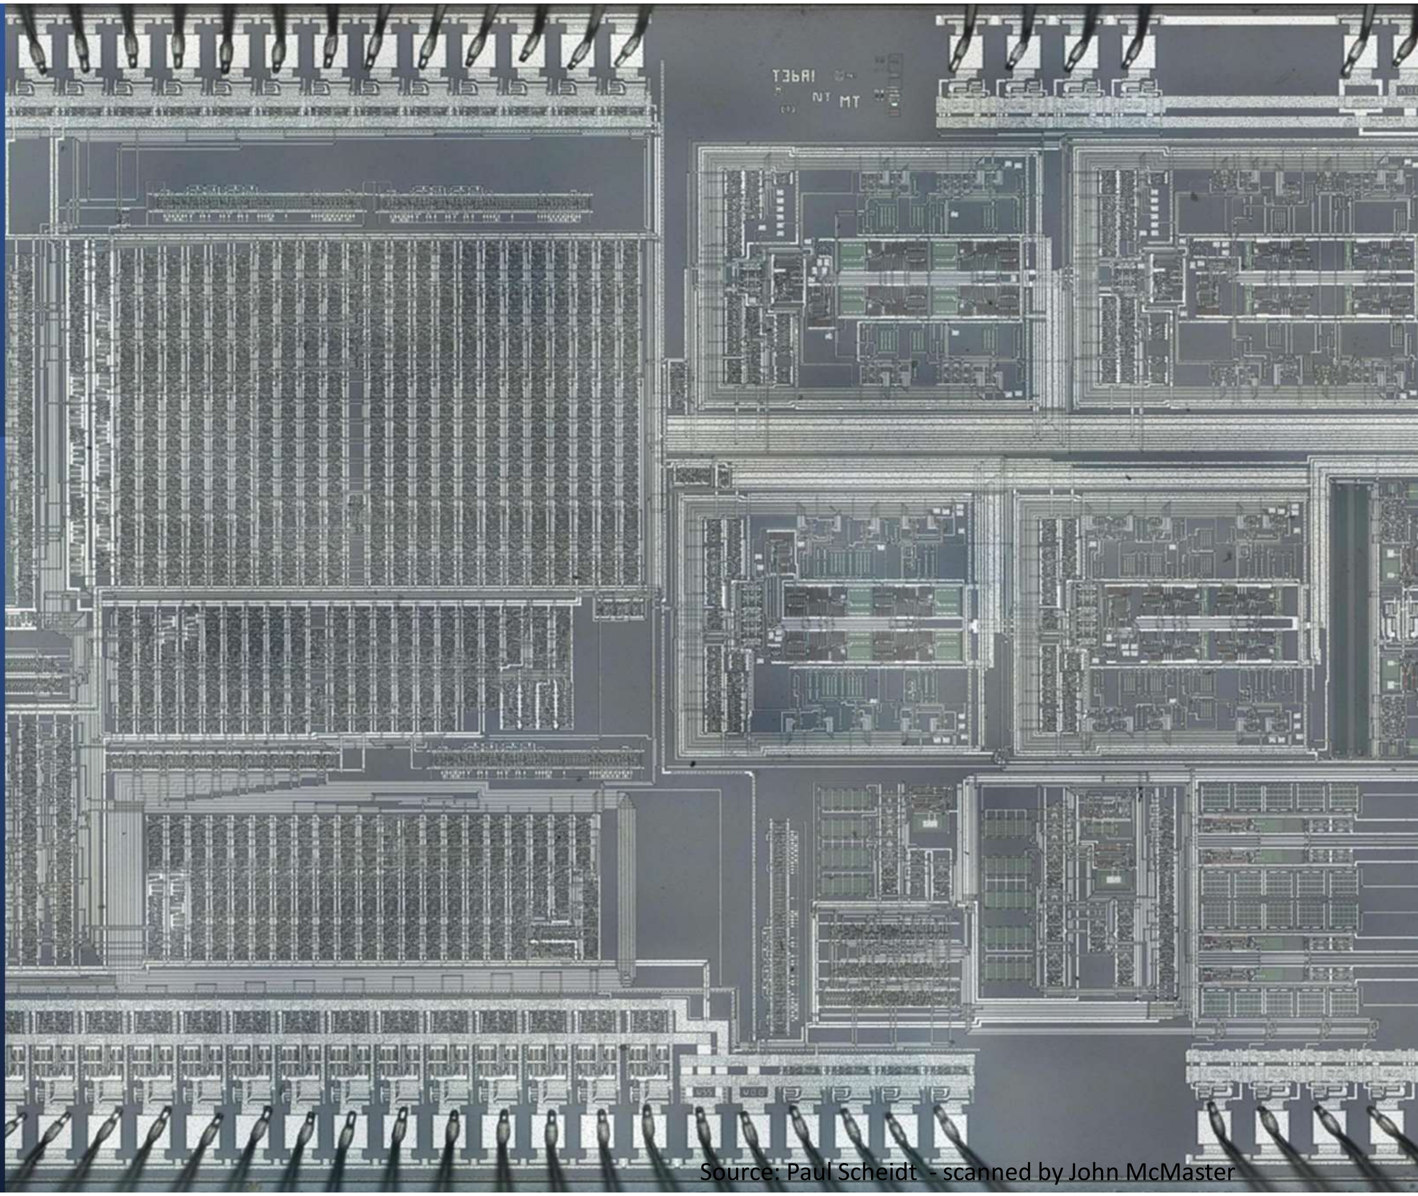

Source: Paul Scheidt - scanned by John McMaster

Motorola MC68000 Microarchitecture

Two Level Microinstruction Control

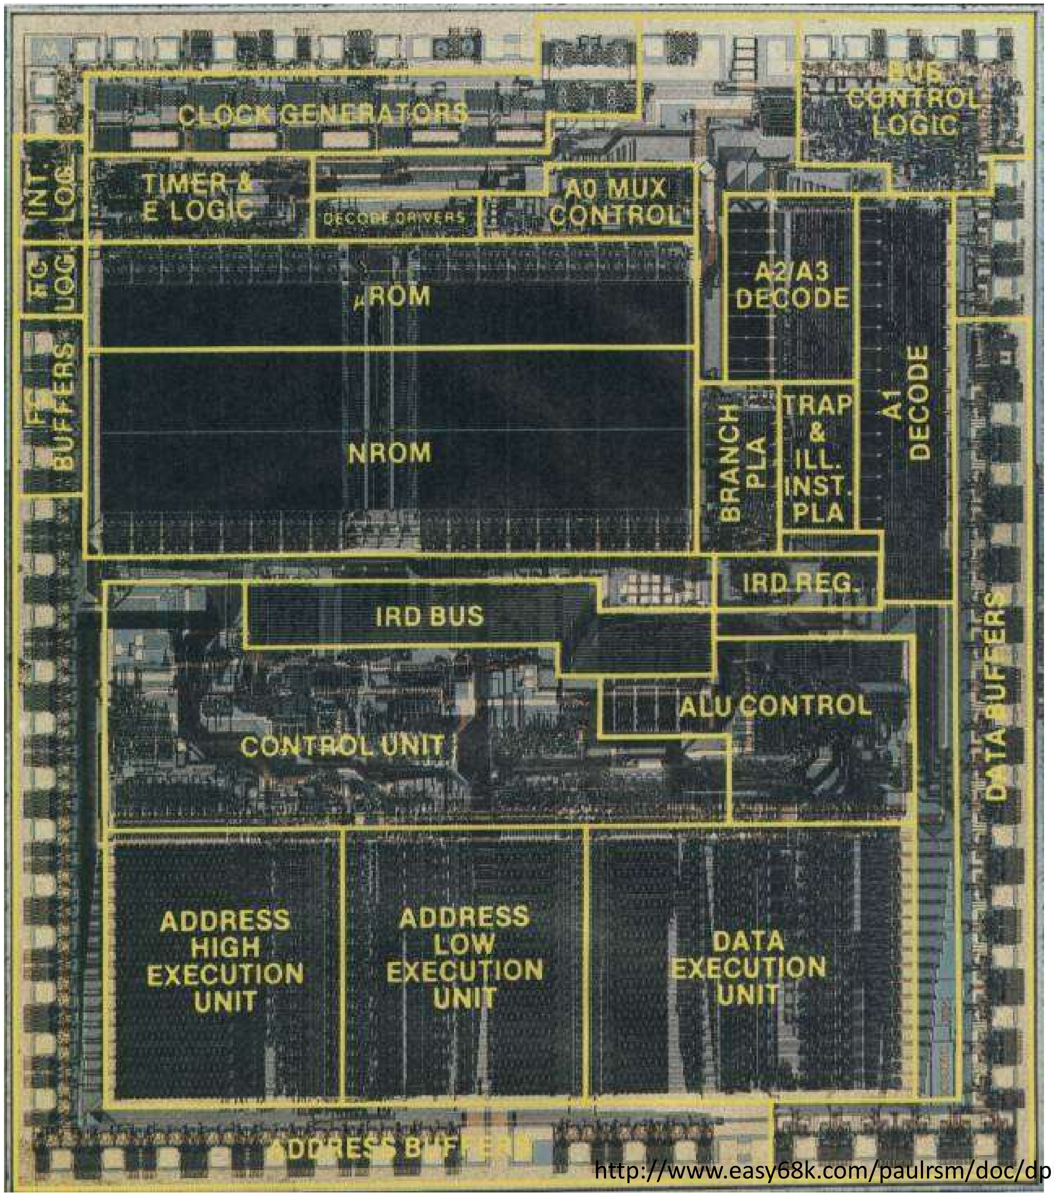

http://www.easy68k.com/paulrsm/doc/dpbm68k2.htm

Two Level Microinstruction Control

Can we do better?

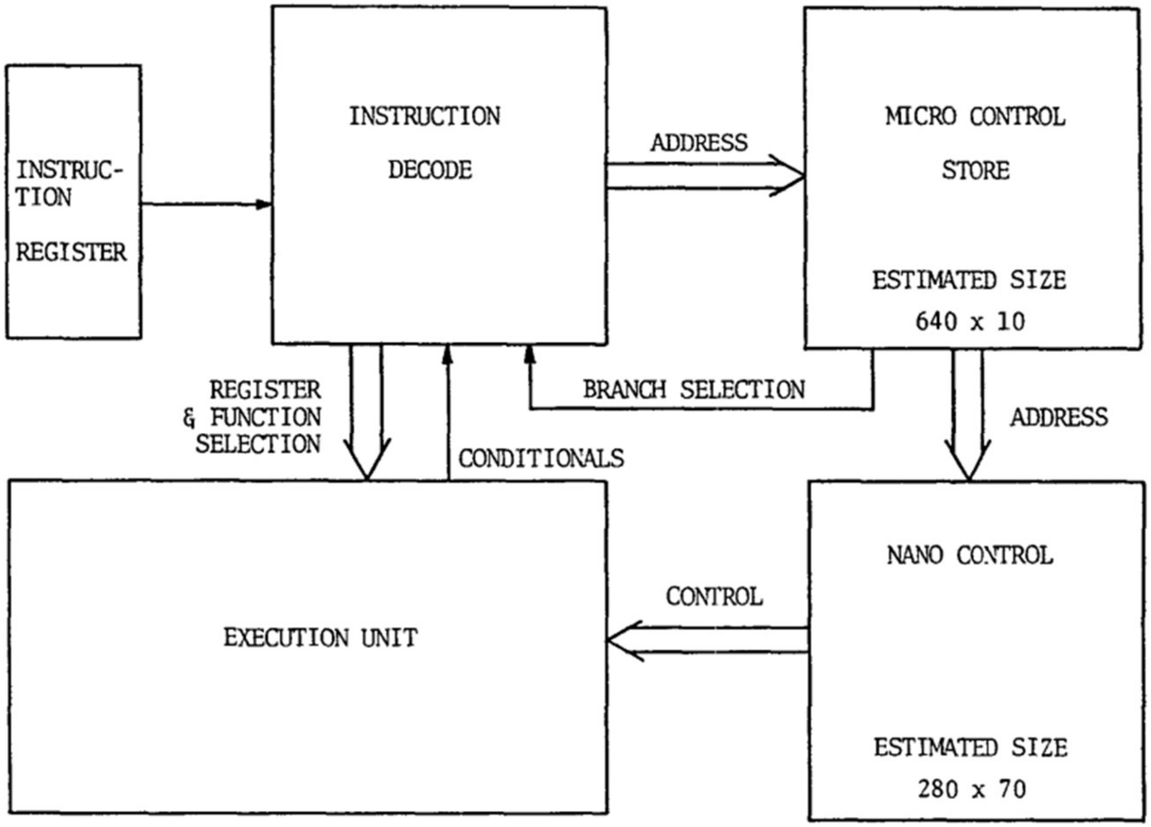

FIGURE 3. MC68000 CONTROL STRUCTURE

MICROPROGRAMMED IMPLEMENTATION OF A SINGLE CHIP MICROPROCESSOR
Skip Stritter & Nick Tredennick, Motorola Semiconductor Group, Austin, Texas

Programmable Mixed-Signal Digital Signal Processor

Bell Northern Research A37

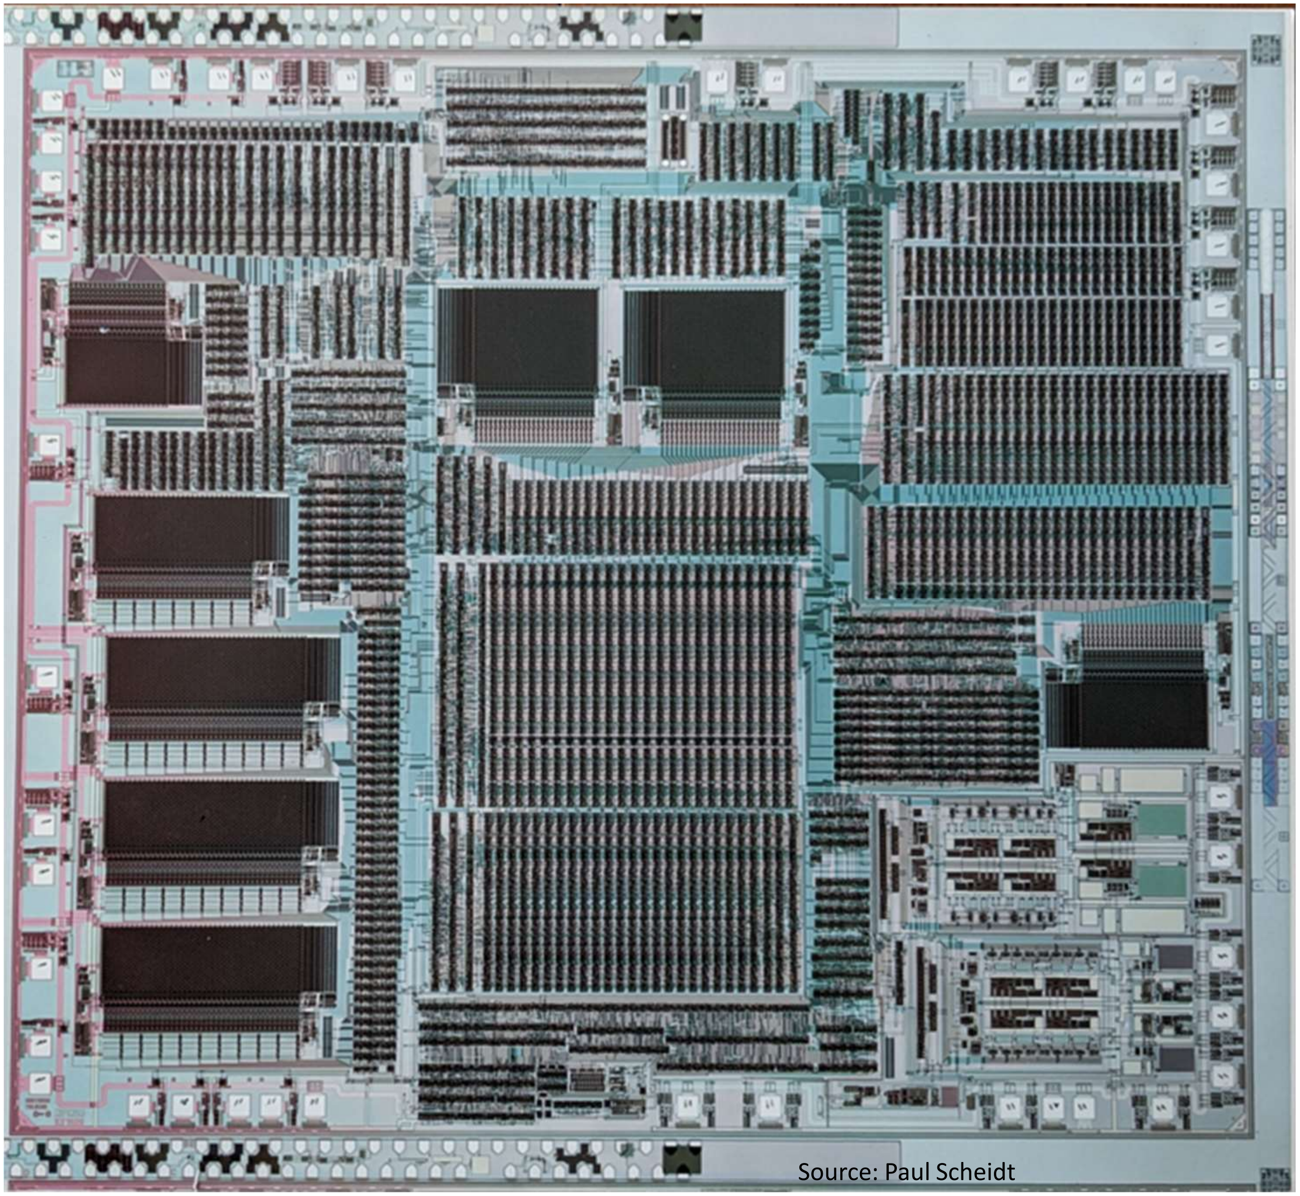

Source: Paul Scheidt

BIST
Built-in
Self Test

Fault Detection in many Scenarios

LFSRs are super Versatile!

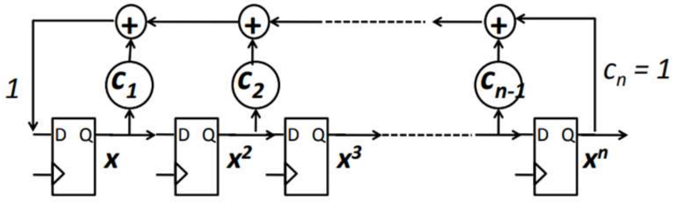

$$P(x) = 1 + c_1\, x + c_2\, x^2 + \ldots + c_n\, x^n$$

LFSR Data Vector Generator

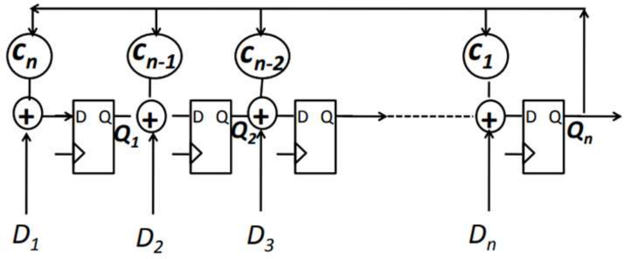

LFSR Signature Compressor

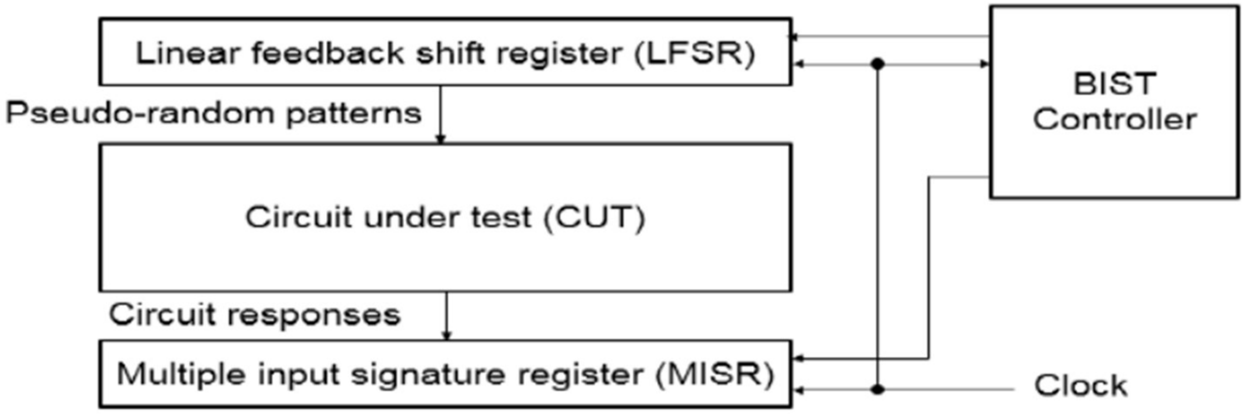

Linear feedback shift register (LFSR)

Pseudo-random patterns

Circuit under test (CUT)

Circuit responses

Multiple input signature register (MISR)

BIST Controller

Clock

http://www.dejazzer.com/ece470/resources/slides16.pdf

Zycad
Gate Simulation
Accelerator

late 80s – mid 90s

Test Vector
Fault Grading

Z01X
Modern Fault Simulation

Evaluate Fault Tolerance

Effectiveness of Fault Attack Countermeasures

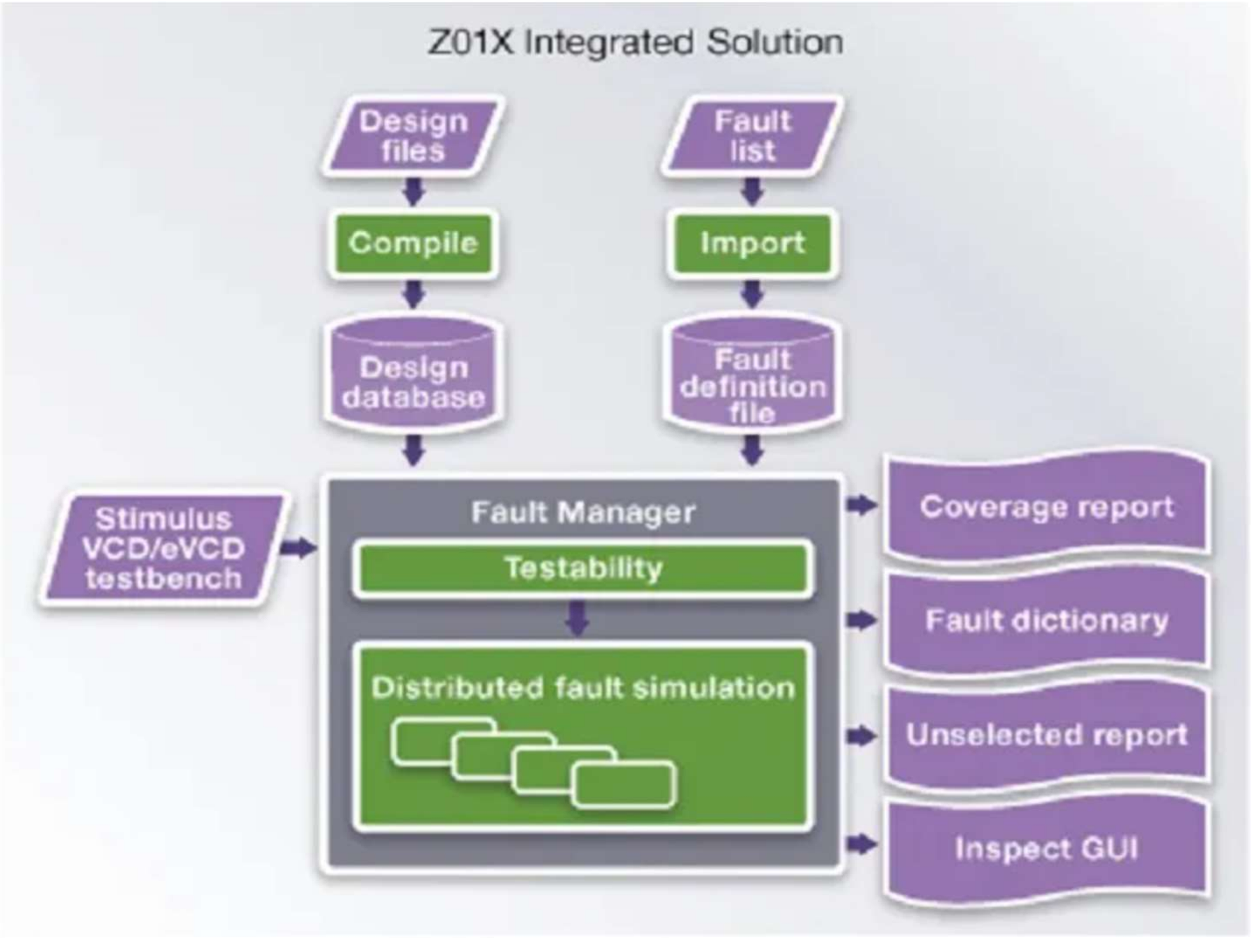

Approaching the 1 micron "Barrier"

Transistors still easily visible

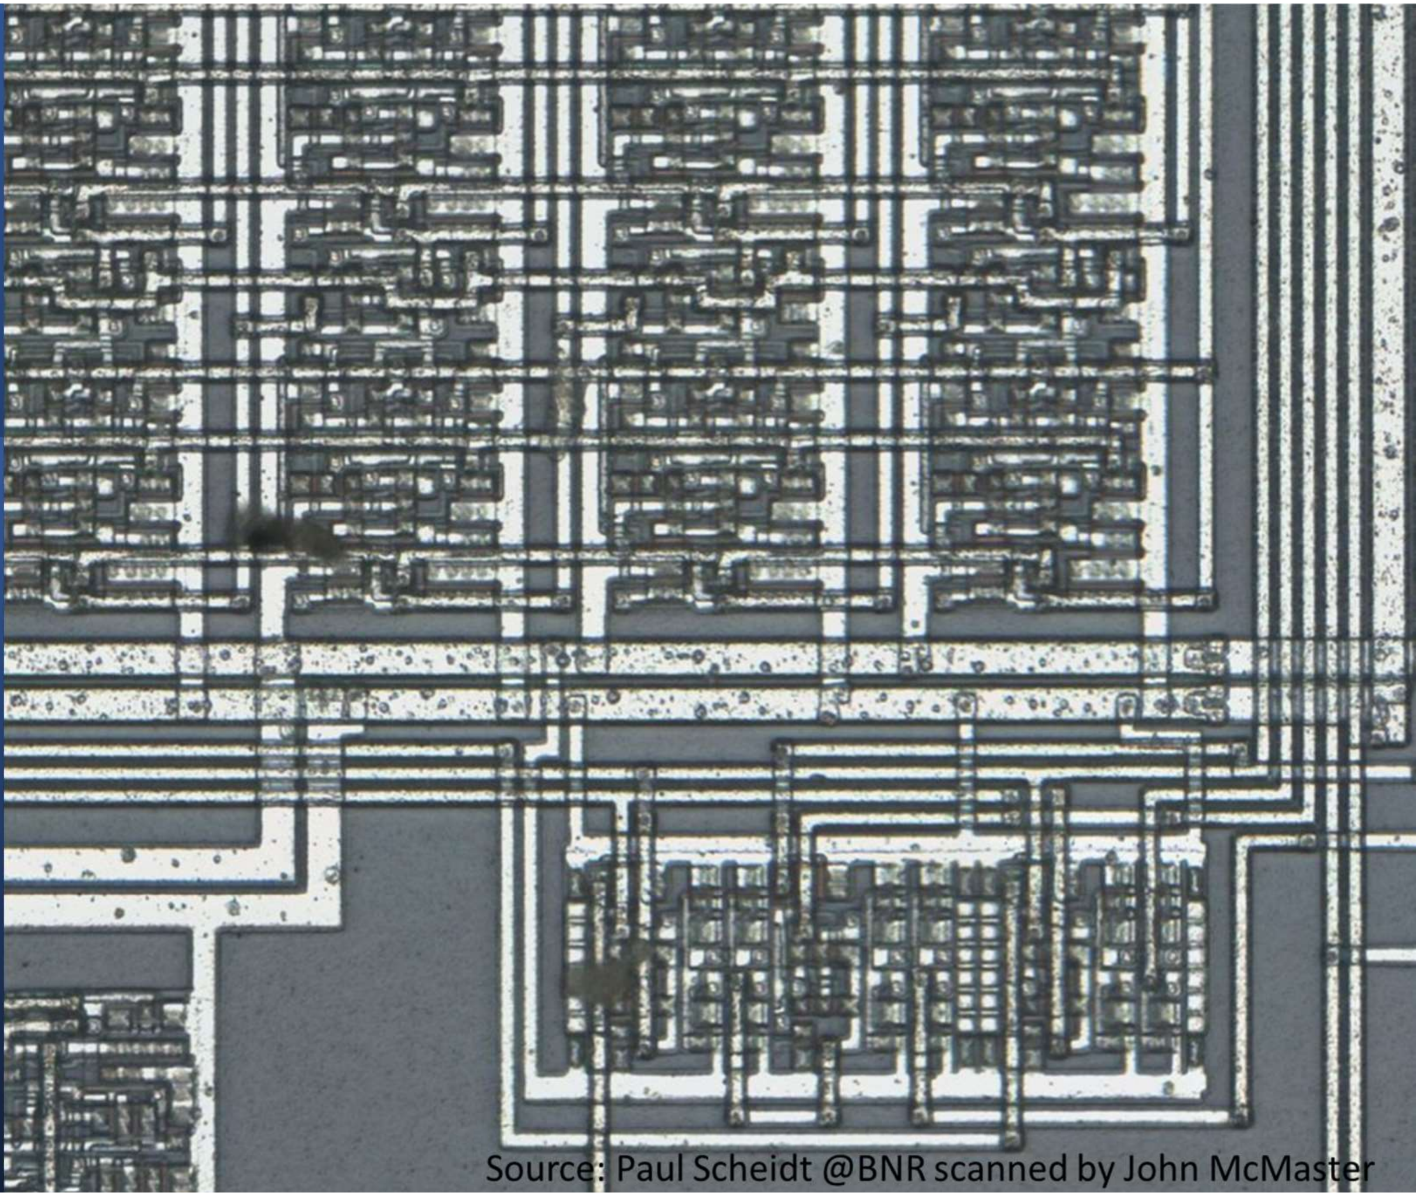

Source: Paul Scheidt @BNR scanned by John McMaster

Limited Metal Interconnect

Access by “drilling” and “wiring”

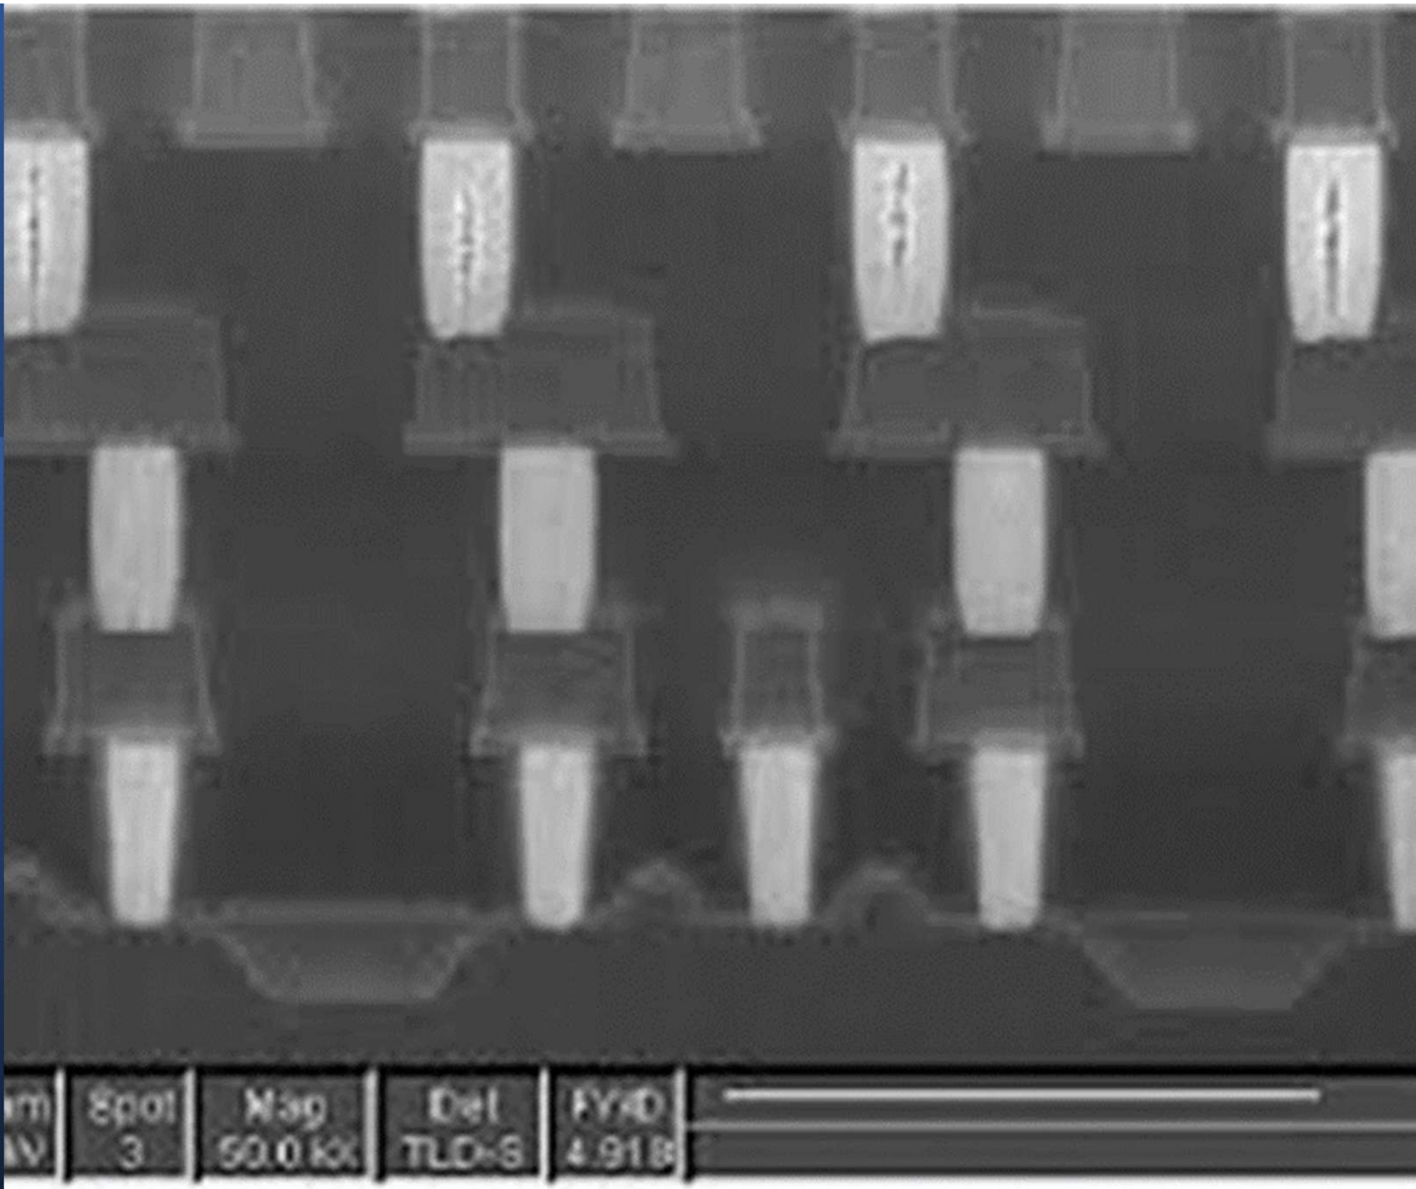

Debugging Silicon with Dynamic Voltage Contrast

Seeing the signals

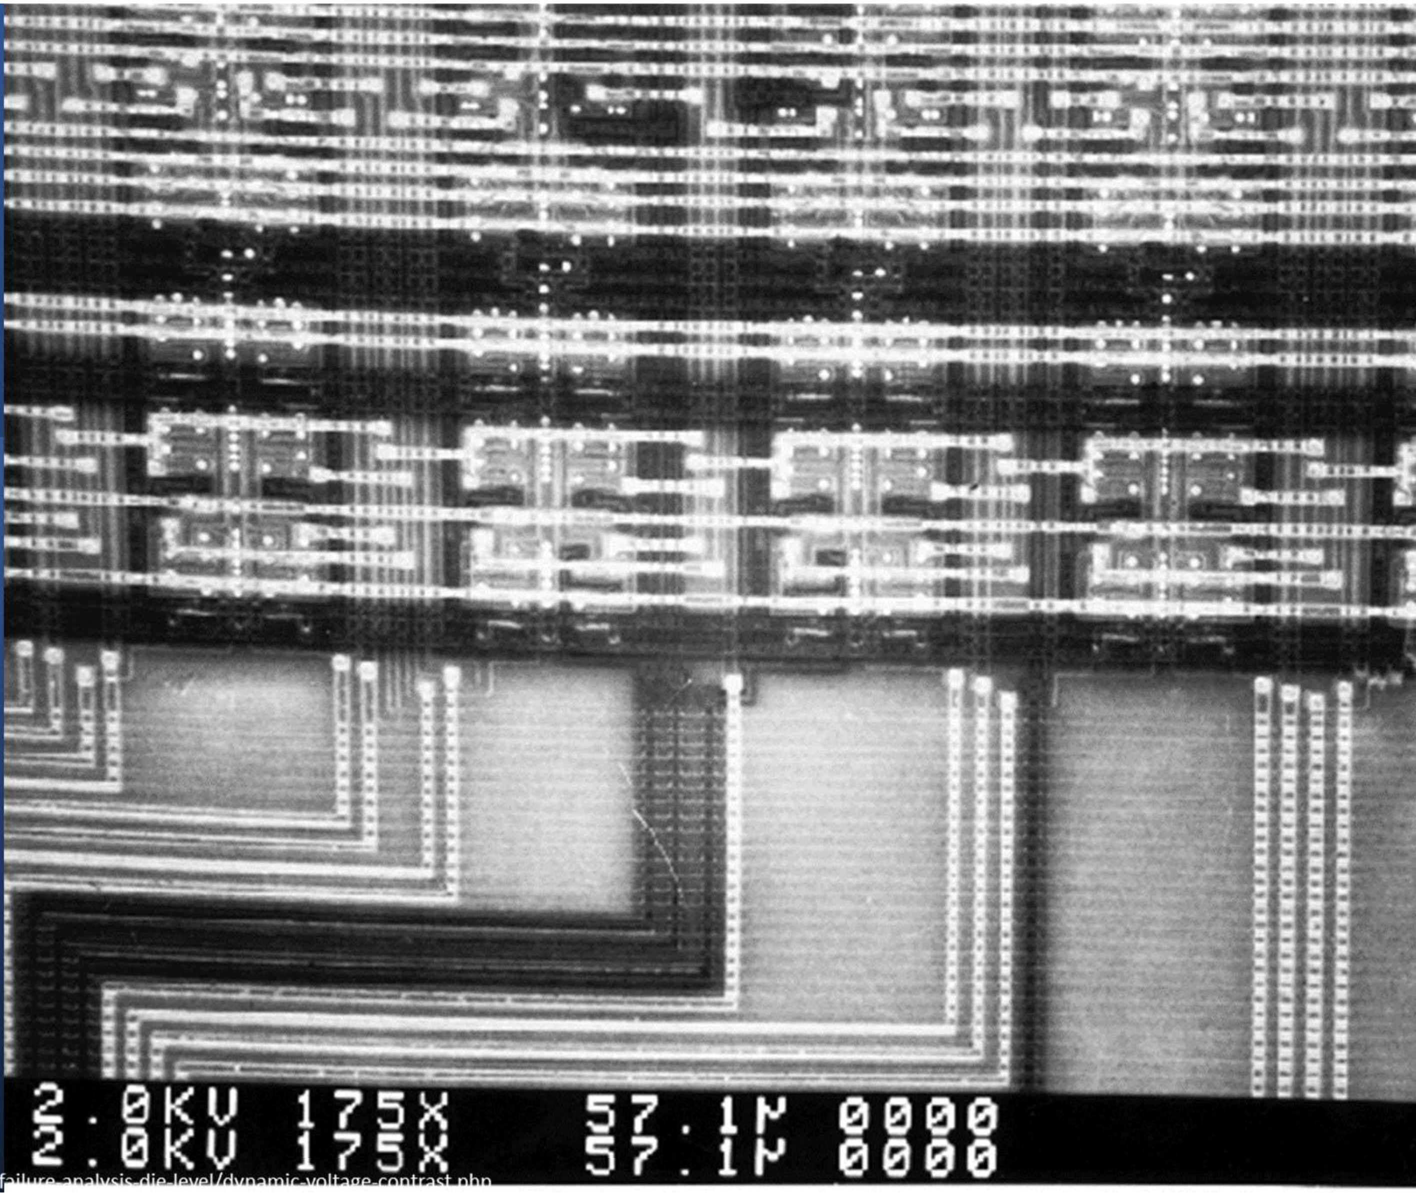

https://www.semitracks.com/reference-material/failure-and-yield-analysis/failure-analysis-die-level/dynamic-voltage-contrast.php

Fixing Silicon with Focused Ion Beam (FIB)

Circuit Editing

https://www.eag.com/services/engineering/fib-circuit-edit-debug/

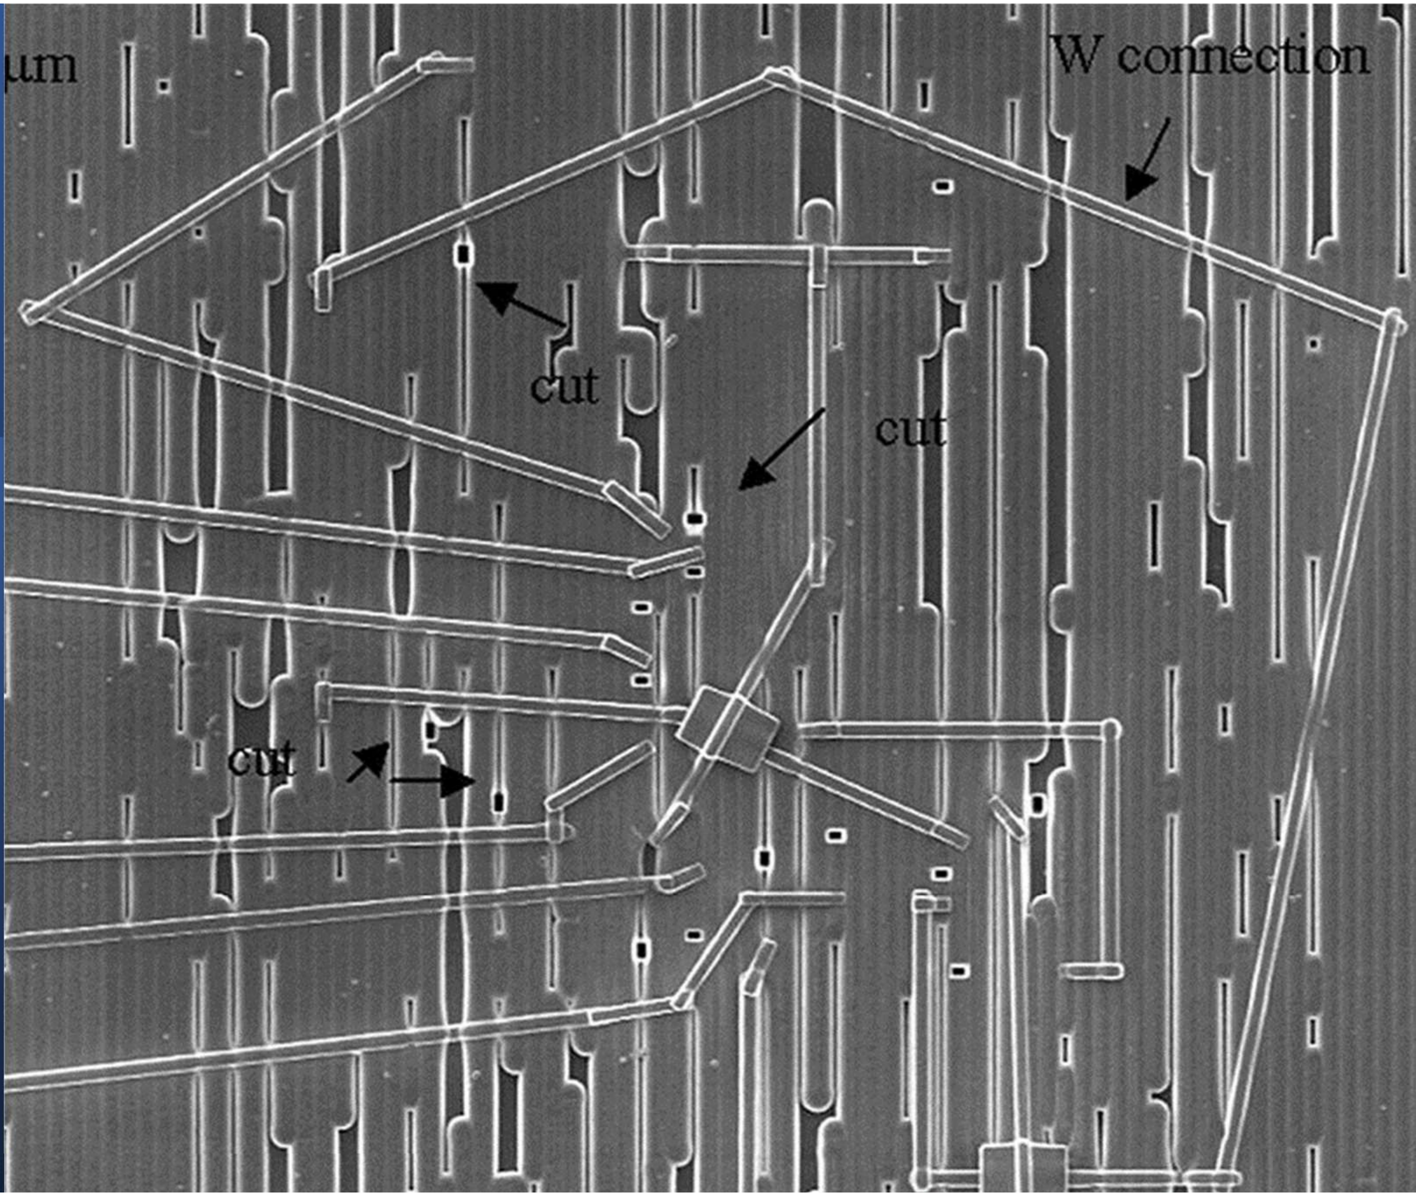

Dealing with increasing Complexity

Abstraction Layers

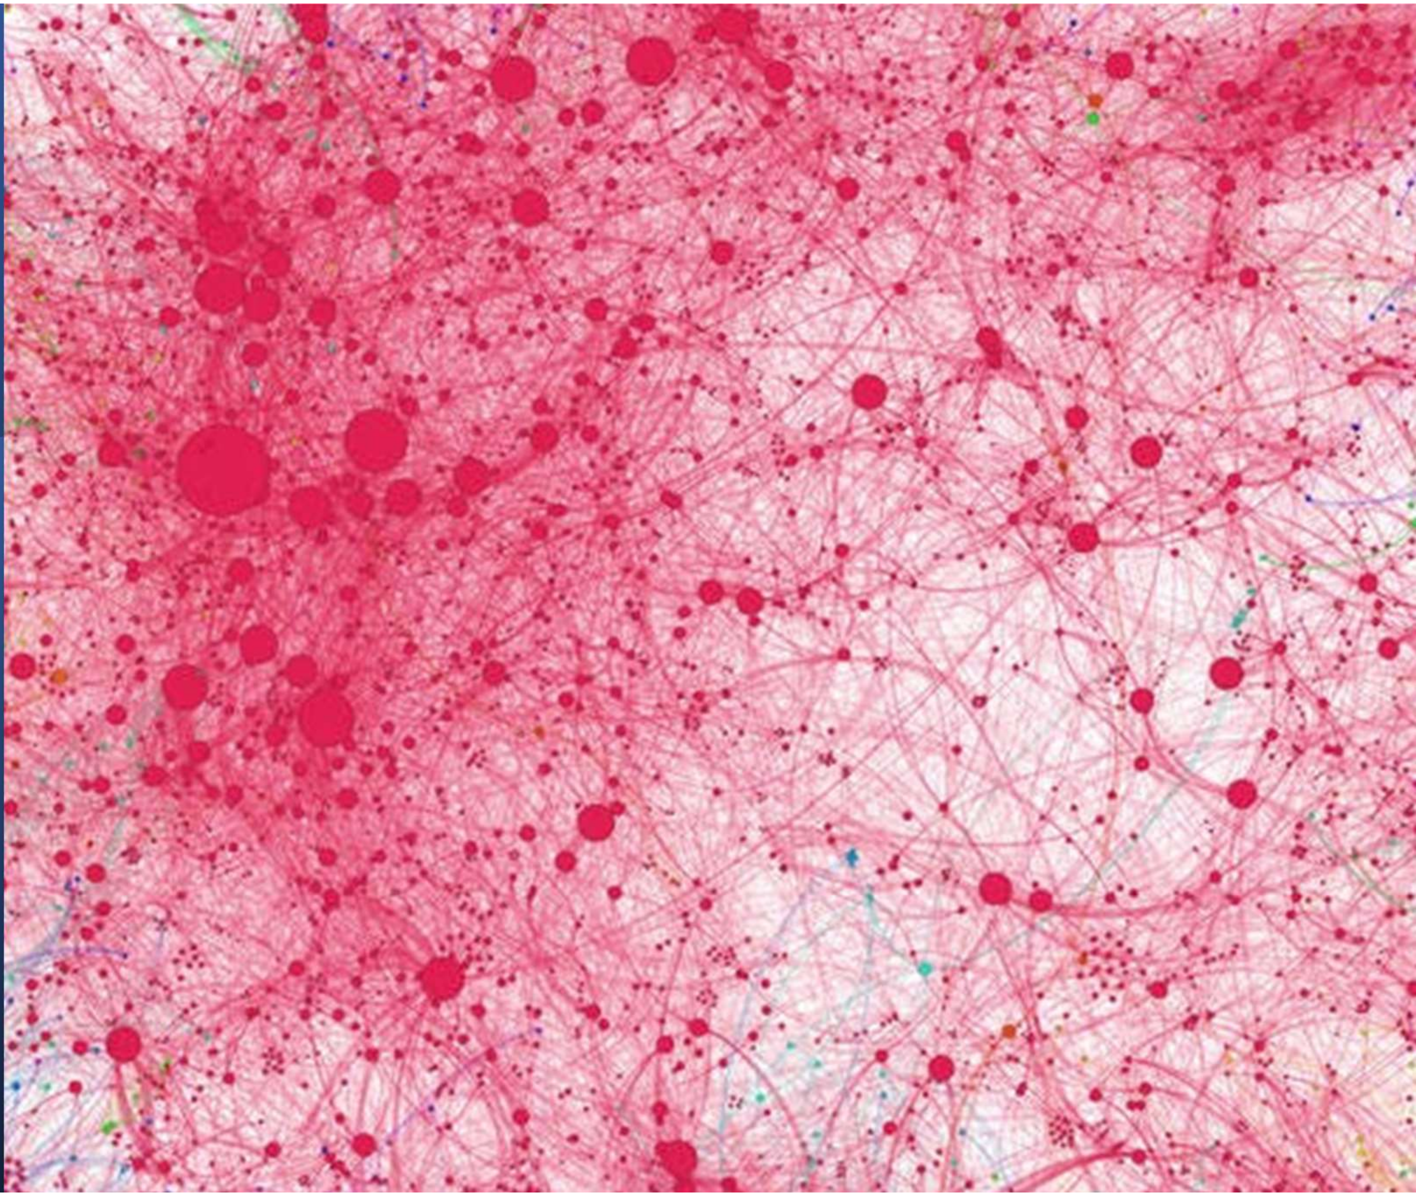

Growing Capacity per Wafer
Shrinking feature sizes 5 to 1 micron

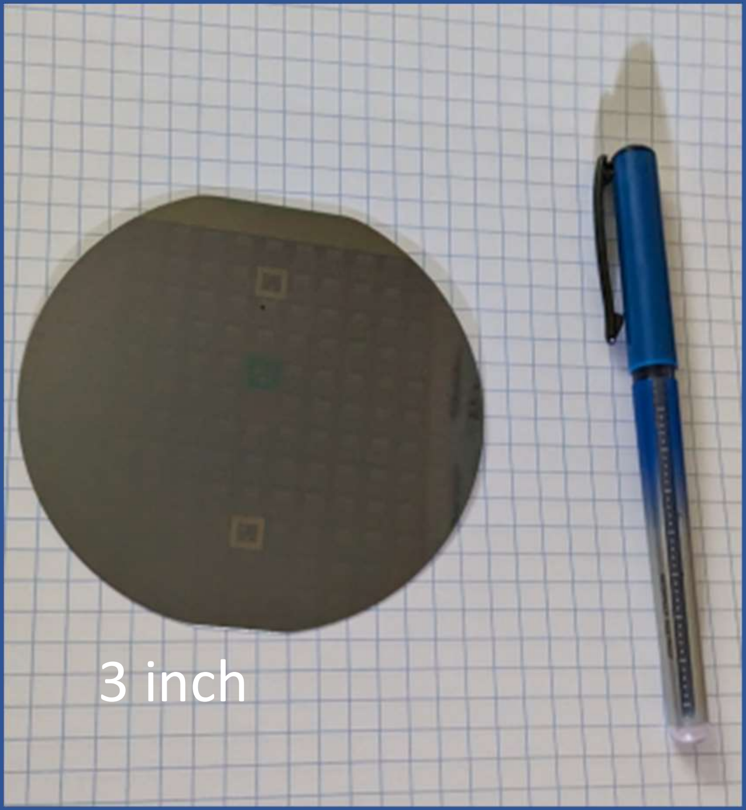

3 inch

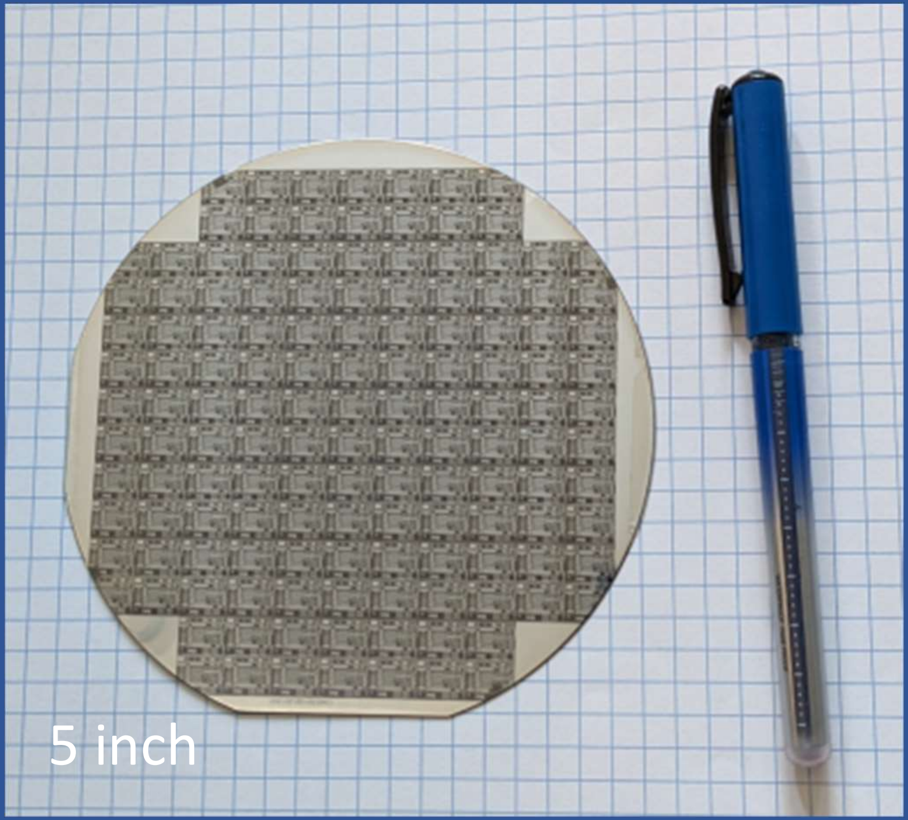

5 inch

Source: Paul Scheidt

Standard Cell Libraries

Structured Custom Design

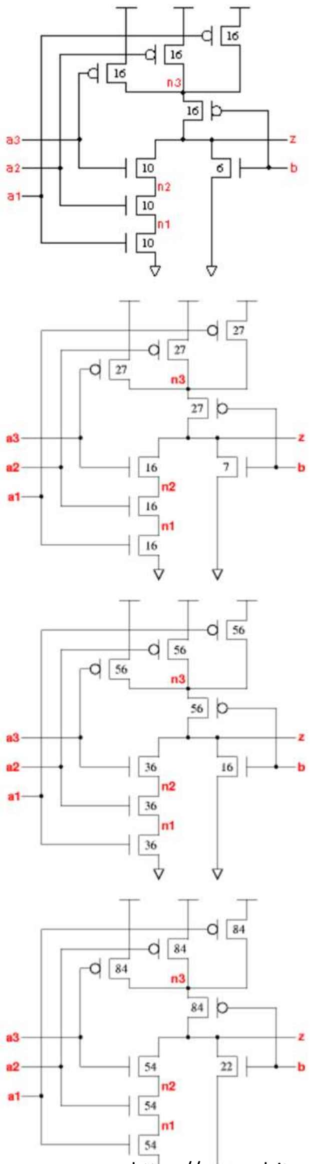

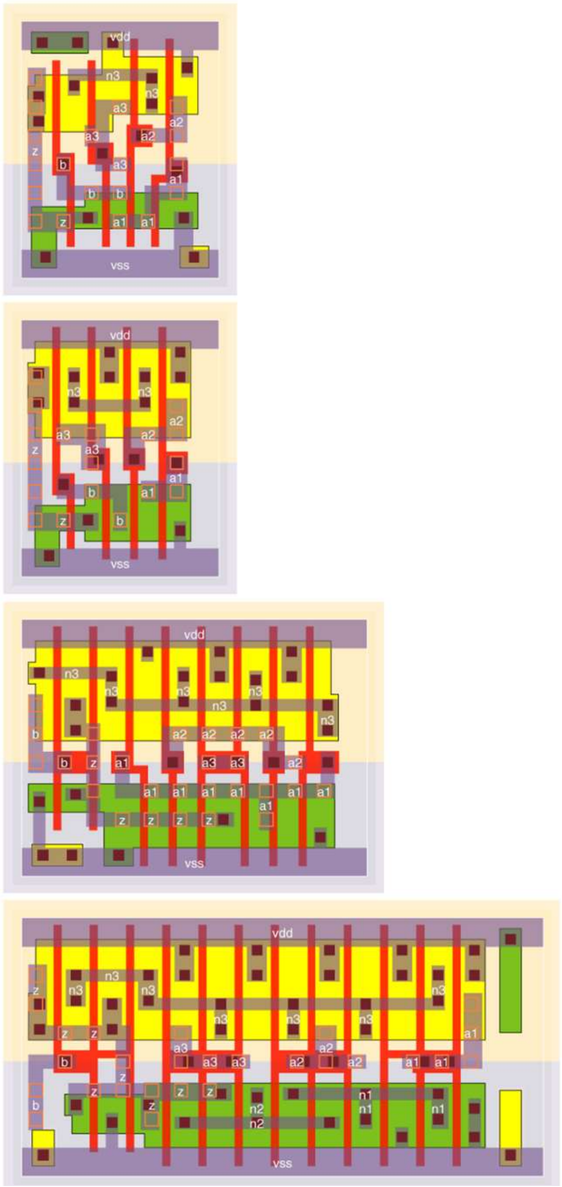

https://www.vlsitechnology.org/html/cells/vsclib013/aoi31.html

Logic Level Design Abstraction

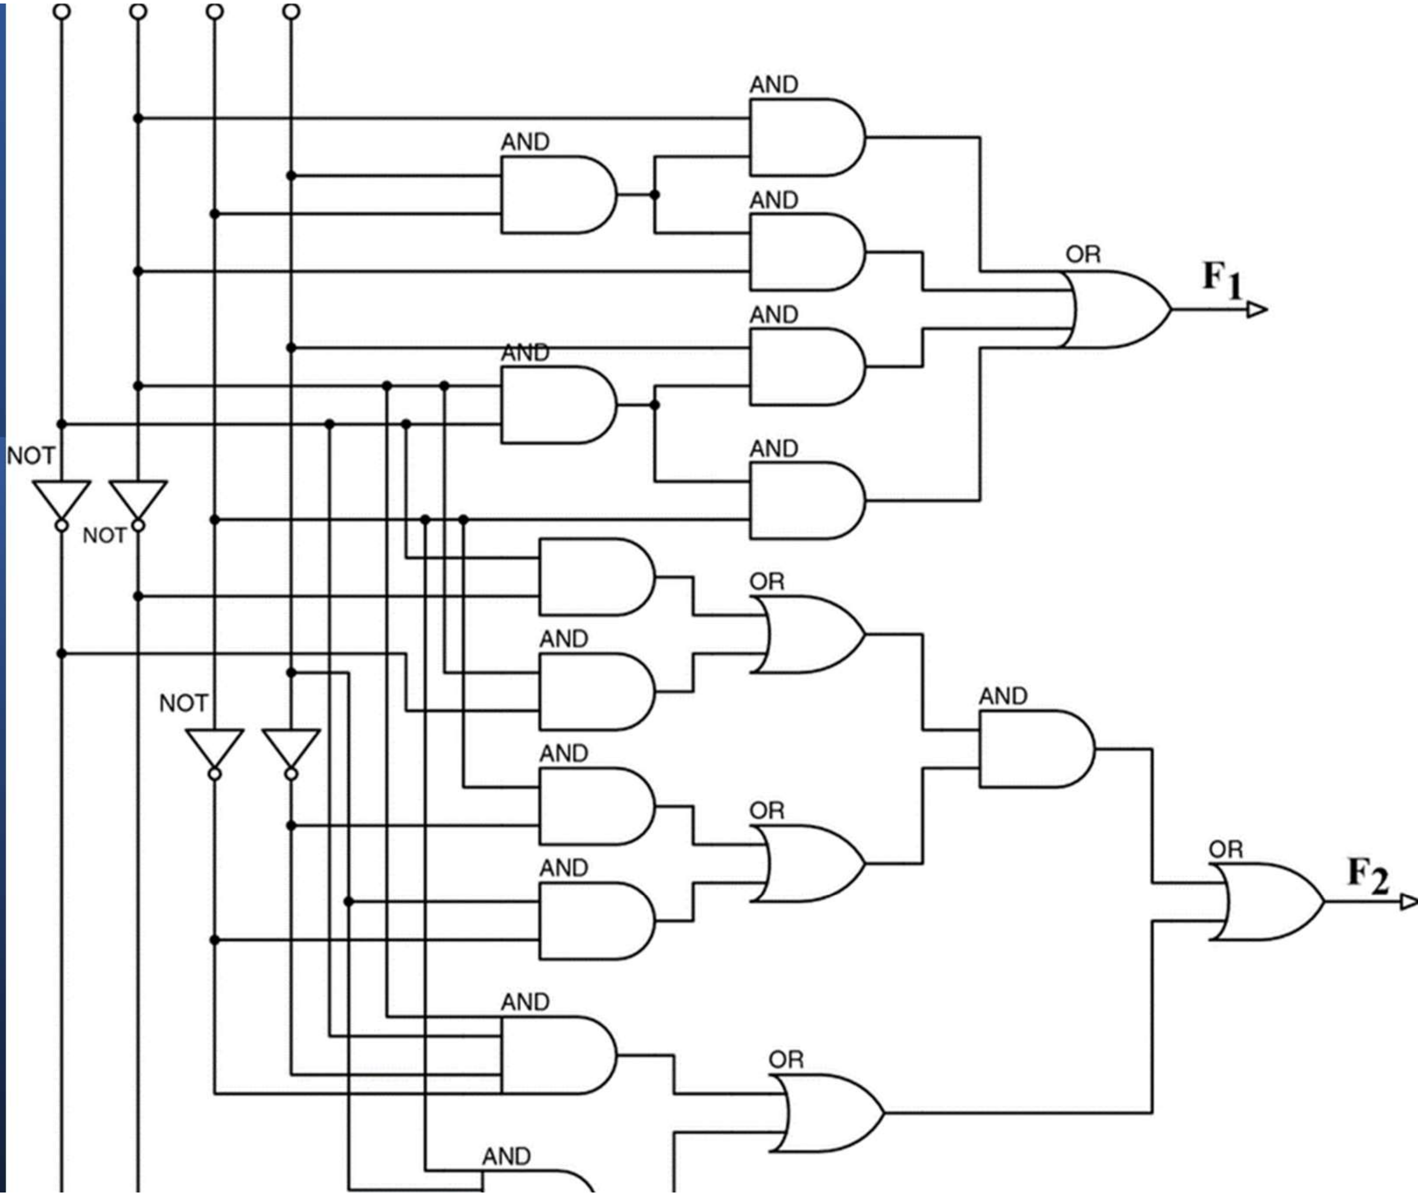

Increasing Abstraction

Logic Gates to Functional

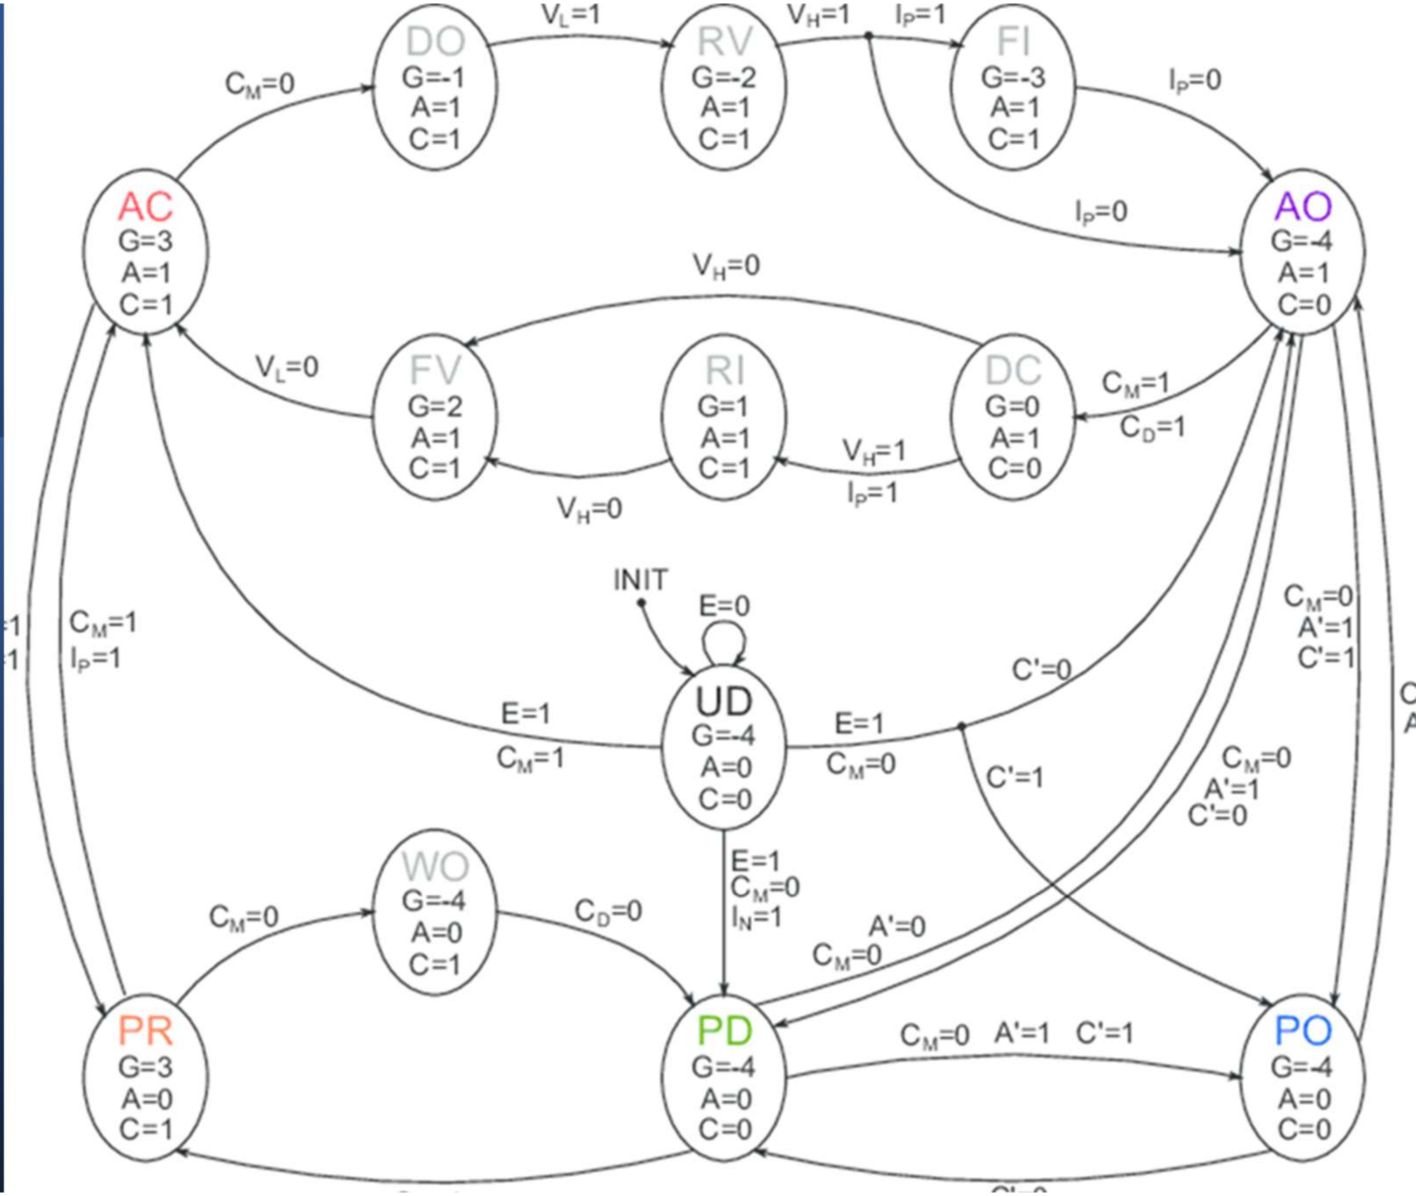

Logic Synthesis and Optimization Revolution

Synopsys Design Compiler

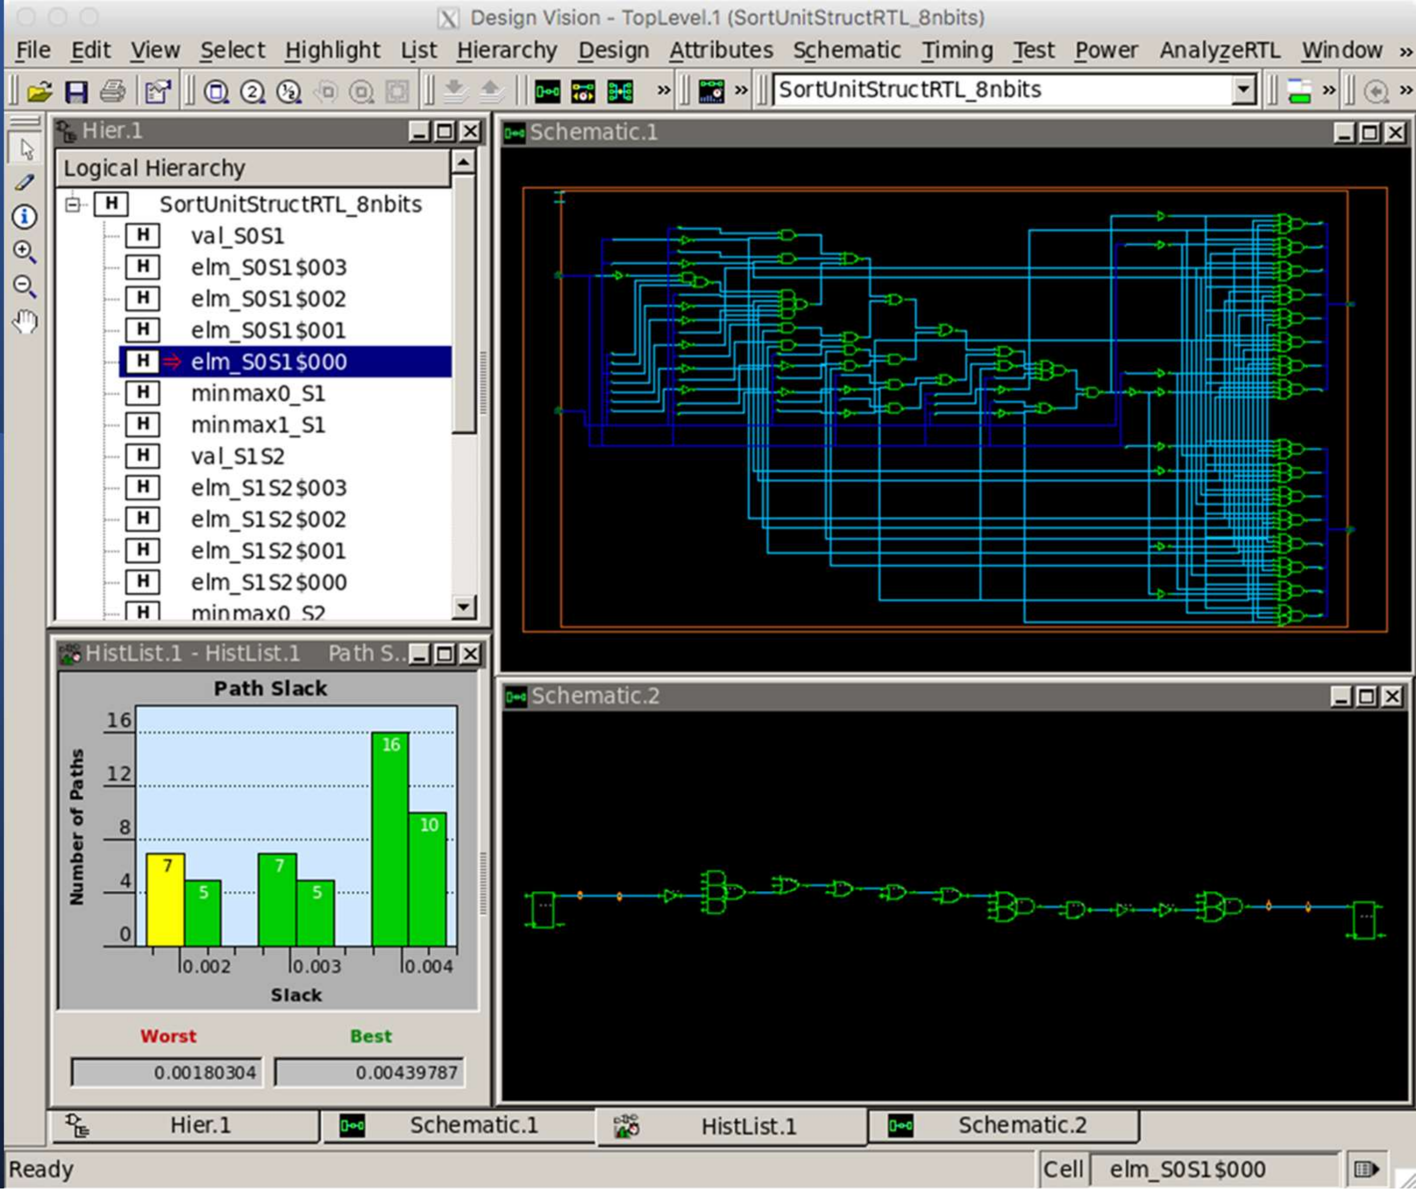

Design Vision - TopLevel.1 (SortUnitStructRTL_8nbits)
File Edit View Select Highlight List Hierarchy Design Attributes Schematic Timing Test Power AnalyzeRTL Window »
SortUnitStructRTL_8nbits
Hier.1
Schematic.1
HistList.1 - HistList.1 Path S...
Schematic.2
Path Slack
Number of Paths
Slack
16
12
8
4
0
0.002
0.003
0.004
7
5
7
5
16
10
Hier.1
Schematic.1
HistList.1
Schematic.2

Logical Hierarchy

- SortUnitStructRTL_8nbits
  - val_S0S1
  - elm_S0S1$003
  - elm_S0S1$002
  - elm_S0S1$001
  - elm_S0S1$000
  - minmax0_S1
  - minmax1_S1
  - val_S1S2
  - elm_S1S2$003
  - elm_S1S2$002
  - elm_S1S2$001
  - elm_S1S2$000
  - minmax0_S2

Worst	Best
0.00180304	0.00439787

Ready

Cell elm_S0S1$000

EDA Industry takes off in the 1990s

Customer In-house Tools mostly replaced

CALMA

daisy
SYSTEMS CORP.

Mentor Graphics®

GATEWAY
DESIGN

VALID

cādence™

VIEWlogic®

SYNOPSYS®

High Level Language Silicon Compilation

VHDL & Verilog

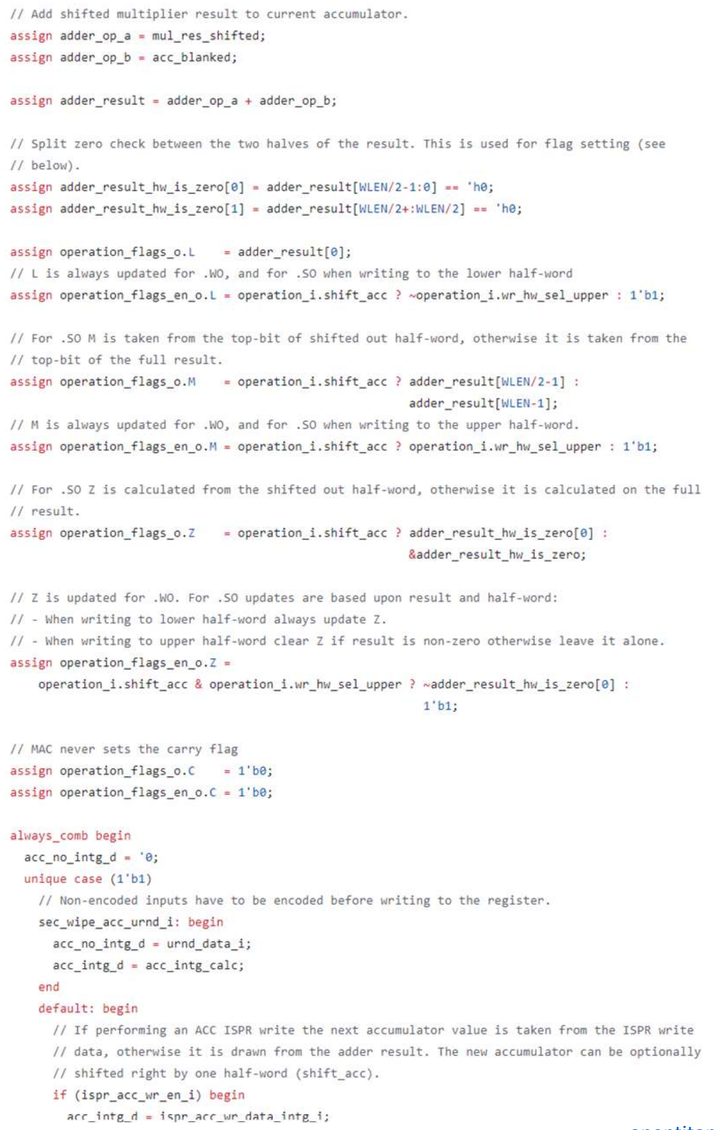

```
// Add shifted multiplier result to current accumulator.
assign adder_op_a = mul_res_shifted;
assign adder_op_b = acc_blanked;

assign adder_result = adder_op_a + adder_op_b;

// Split zero check between the two halves of the result. This is used for flag setting (see
// below).
assign adder_result_hw_is_zero[0] = adder_result[WLEN/2-1:0] == 'h0;
assign adder_result_hw_is_zero[1] = adder_result[WLEN/2+:WLEN/2] == 'h0;

assign operation_flags_o.L    = adder_result[0];
// L is always updated for .WO, and for .SO when writing to the lower half-word
assign operation_flags_en_o.L = operation_i.shift_acc ? ~operation_i.wr_hw_sel_upper : 1'b1;

// For .SO M is taken from the top-bit of shifted out half-word, otherwise it is taken from the
// top-bit of the full result.
assign operation_flags_o.M    = operation_i.shift_acc ? adder_result[WLEN/2-1] :
                                                        adder_result[WLEN-1];
// M is always updated for .WO, and for .SO when writing to the upper half-word.
assign operation_flags_en_o.M = operation_i.shift_acc ? operation_i.wr_hw_sel_upper : 1'b1;

// For .SO Z is calculated from the shifted out half-word, otherwise it is calculated on the full
// result.
assign operation_flags_o.Z    = operation_i.shift_acc ? adder_result_hw_is_zero[0] :
                                                        &adder_result_hw_is_zero;

// Z is updated for .WO. For .SO updates are based upon result and half-word:
// - When writing to lower half-word always update Z.
// - When writing to upper half-word clear Z if result is non-zero otherwise leave it alone.
assign operation_flags_en_o.Z =
    operation_i.shift_acc & operation_i.wr_hw_sel_upper ? ~adder_result_hw_is_zero[0] :
                                                          1'b1;

// MAC never sets the carry flag
assign operation_flags_o.C    = 1'b0;
assign operation_flags_en_o.C = 1'b0;

always_comb begin
  acc_no_intg_d = '0;
  unique case (1'b1)
    // Non-encoded inputs have to be encoded before writing to the register.
    sec_wipe_acc_urnd_i: begin
      acc_no_intg_d = urnd_data_i;
      acc_intg_d = acc_intg_calc;
    end
    default: begin
      // If performing an ACC ISPR write the next accumulator value is taken from the ISPR write
      // data, otherwise it is drawn from the adder result. The new accumulator can be optionally
      // shifted right by one half-word (shift_acc).
      if (ispr_acc_wr_en_i) begin
        acc_intg_d = ispr_acc_wr_data_intg_i;
```

opentitan/hw/ip/otbn/rtl at master · lowRISC/opentitan · GitHub

Moore’s Law Metrics over 4 Decades

Min Feature Size
1978: 5 µm
2020: 5 nm
1000 : 1

SRAM Cell Density
43,000 : 1

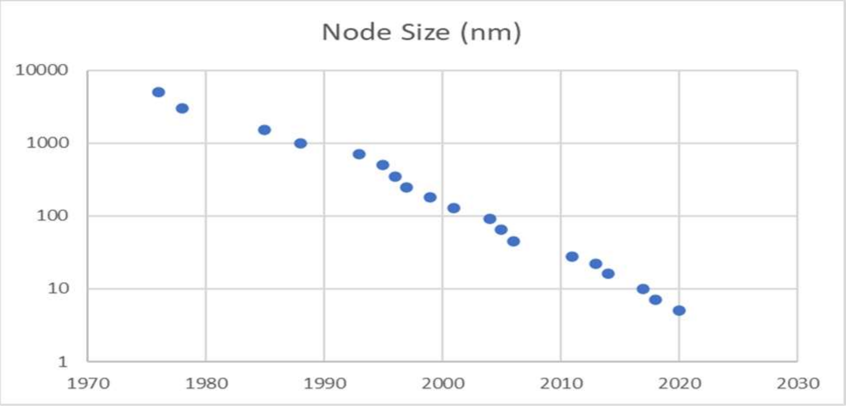
Node Size (nm)

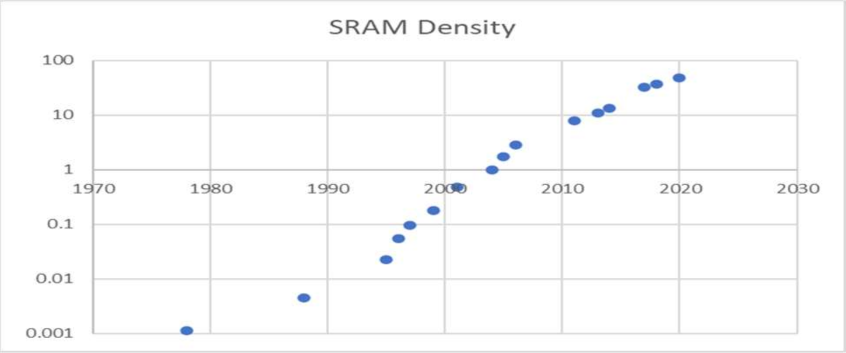
SRAM Density

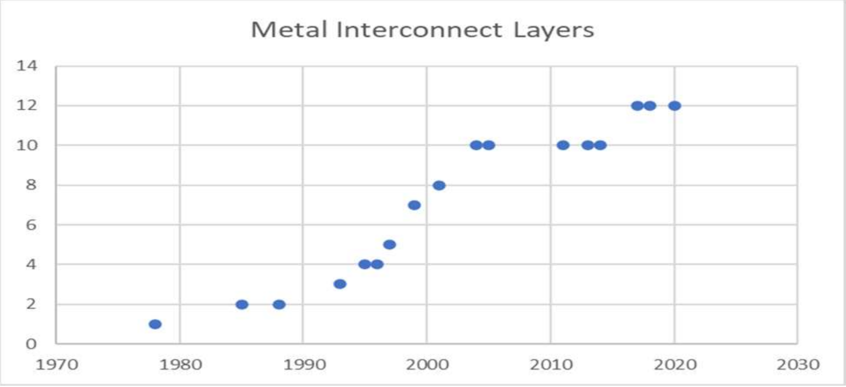
Metal Interconnect Layers

Moore's Law Metrics over 4 Decades

Transistor Count
1980: 20k
2020: 80B
4 Million : 1

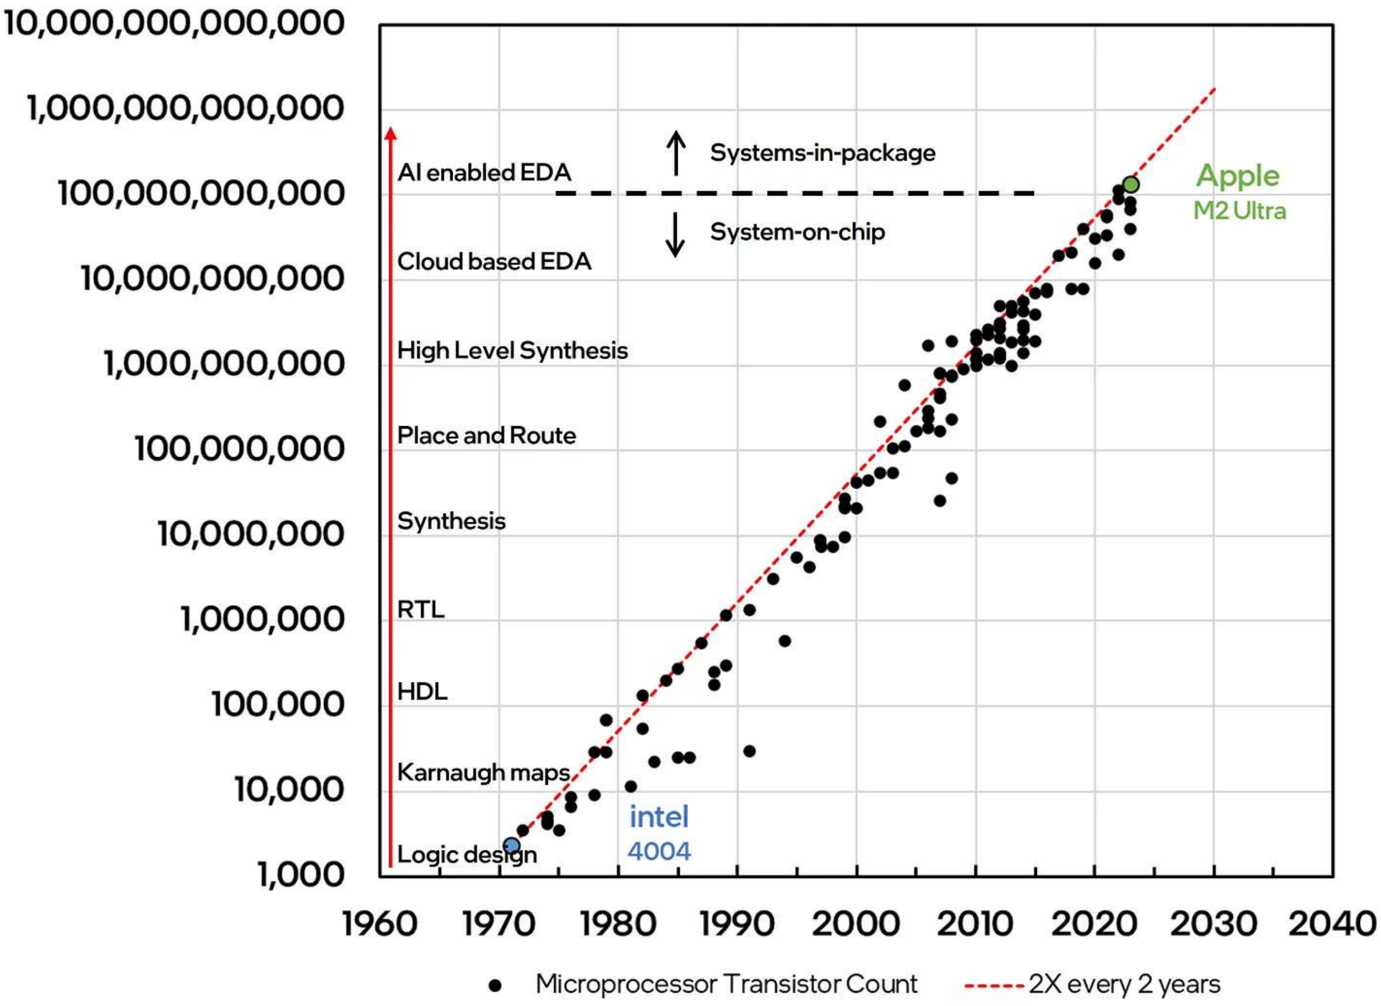

https://semiconductor.substack.com/p/the-relentless-pursuit-of-moores

Wafer Densities Grow Further to 300 mm

Increasing Chip Die Sizes and Production Capacities

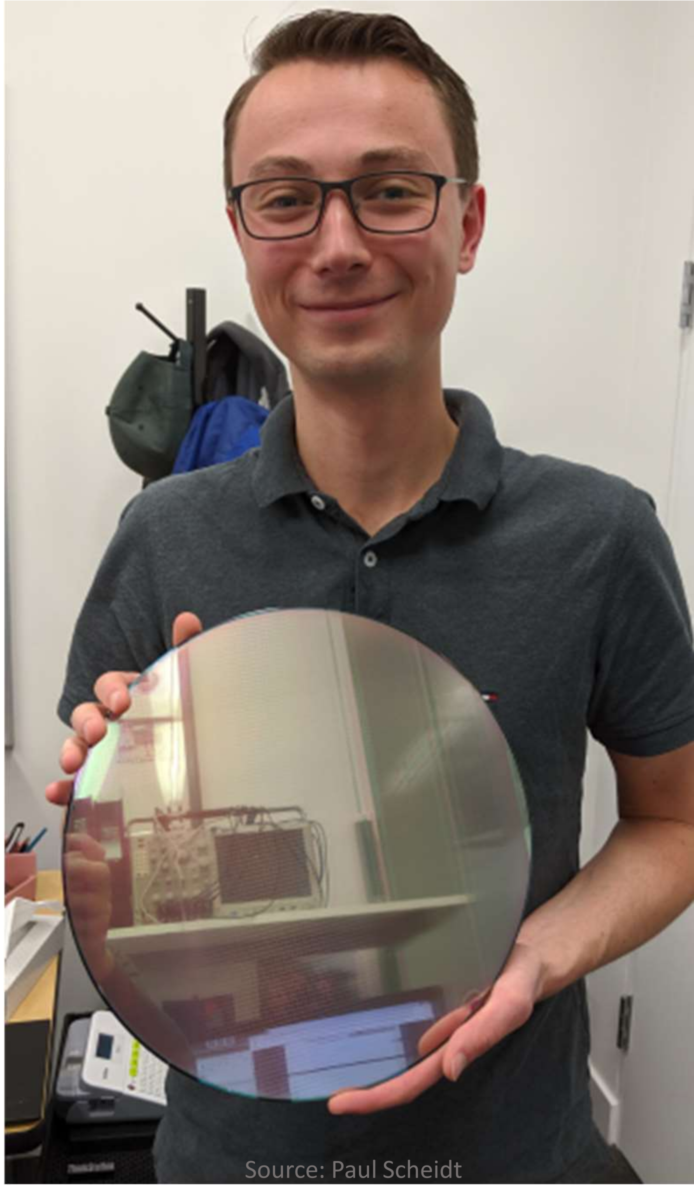

Source: Paul Scheidt

Interconnect Density and FinFET

8+ layers at 22nm

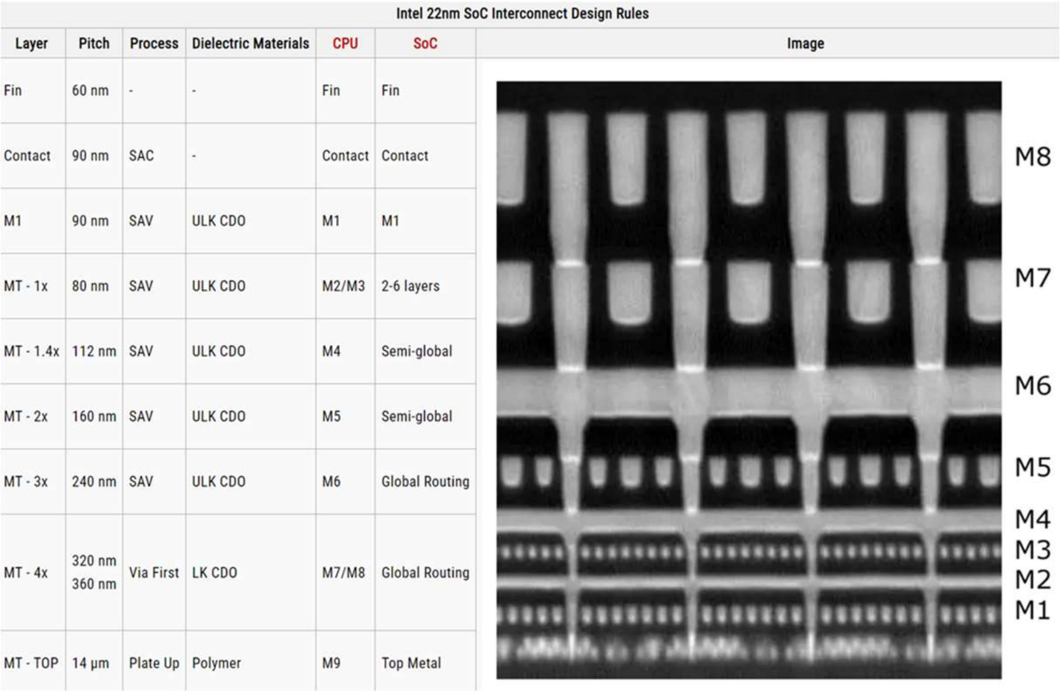

Intel 22nm SoC Interconnect Design Rules

Layer	Pitch	Process	Dielectric Materials	CPU	SoC	Image
Fin	60 nm	-	-	Fin	Fin	
Contact	90 nm	SAC	-	Contact	Contact	
M1	90 nm	SAV	ULK CDO	M1	M1	
MT - 1x	80 nm	SAV	ULK CDO	M2/M3	2-6 layers	
MT - 1.4x	112 nm	SAV	ULK CDO	M4	Semi-global	
MT - 2x	160 nm	SAV	ULK CDO	M5	Semi-global	
MT - 3x	240 nm	SAV	ULK CDO	M6	Global Routing	
MT - 4x	320 nm 360 nm	Via First	LK CDO	M7/M8	Global Routing	
MT - TOP	14 μm	Plate Up	Polymer	M9	Top Metal	

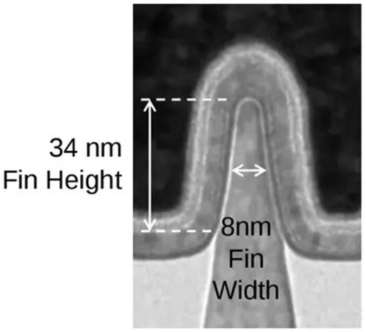

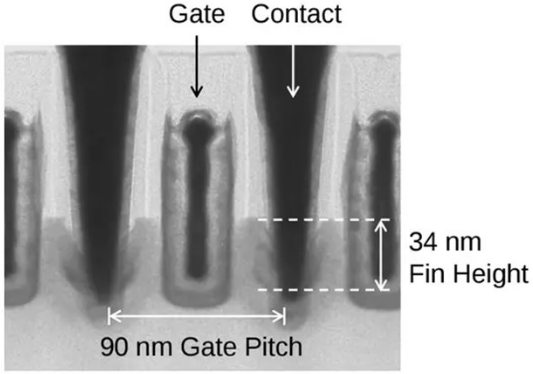

https://en.wikichip.org/wiki/22_nm_lithography_proces

10nm

12+ Metal Layers

Taller & Narrower FinFETS

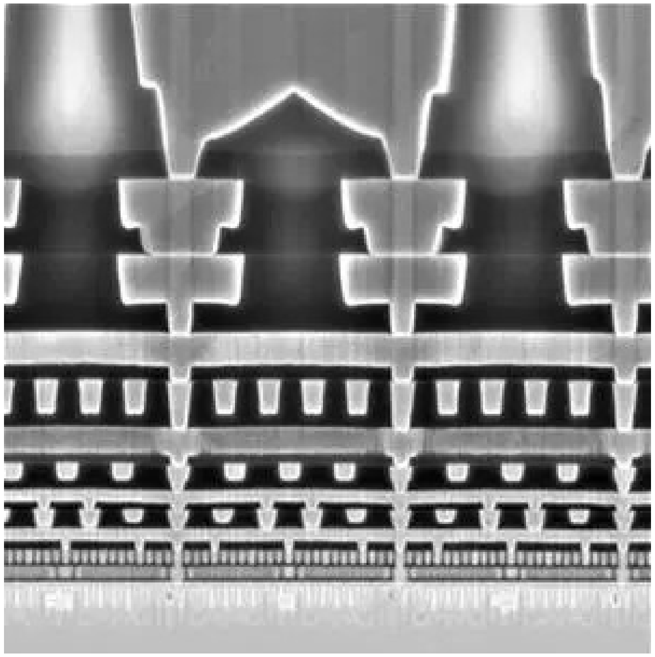

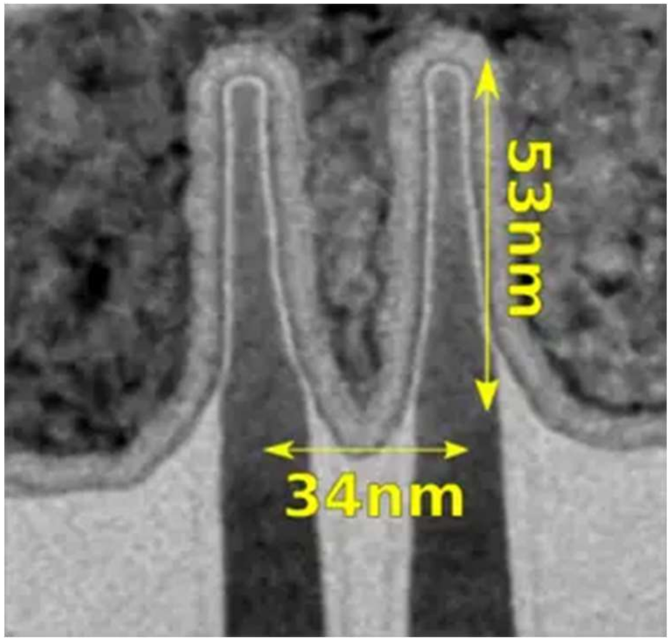

https://en.wikichip.org/wiki/10_nm_lithography_process

Exploding Design Complexity!

Following footsteps of the software industry

opentitan/hw/ip/otbn/rtl at master · lowRISC/opentitan · GitHub

Increase Design Abstraction

Behavioral Compiler

The Technology of the Future?

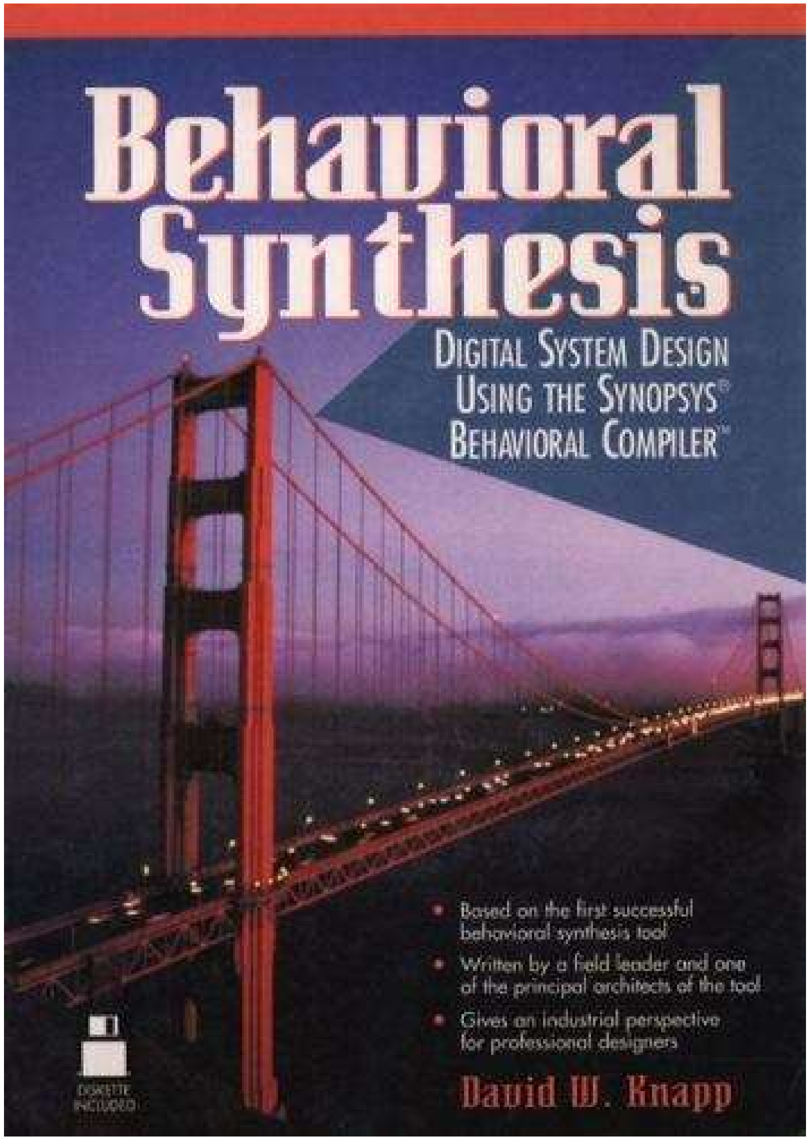

Behavioral Synthesis

DIGITAL SYSTEM DESIGN USING THE SYNOPSYS® BEHAVIORAL COMPILER™

- Based on the first successful behavioral synthesis tool
- Written by a field leader and one of the principal architects of the tool
- Gives an industrial perspective for professional designers

DISKETTE INCLUDED

David W. Knapp

Domain Specific Behavioral Synthesis

More targeted…
More success

tensilica

ARC
Synopsys

Codasip

SYNOPSYS® DESIGNWARE IP DATASHEET

ASIP Designer: Design Tool for Application-Specific Instruction-Set Processors

XLS

Silicon Components

Synopsys DesignWare

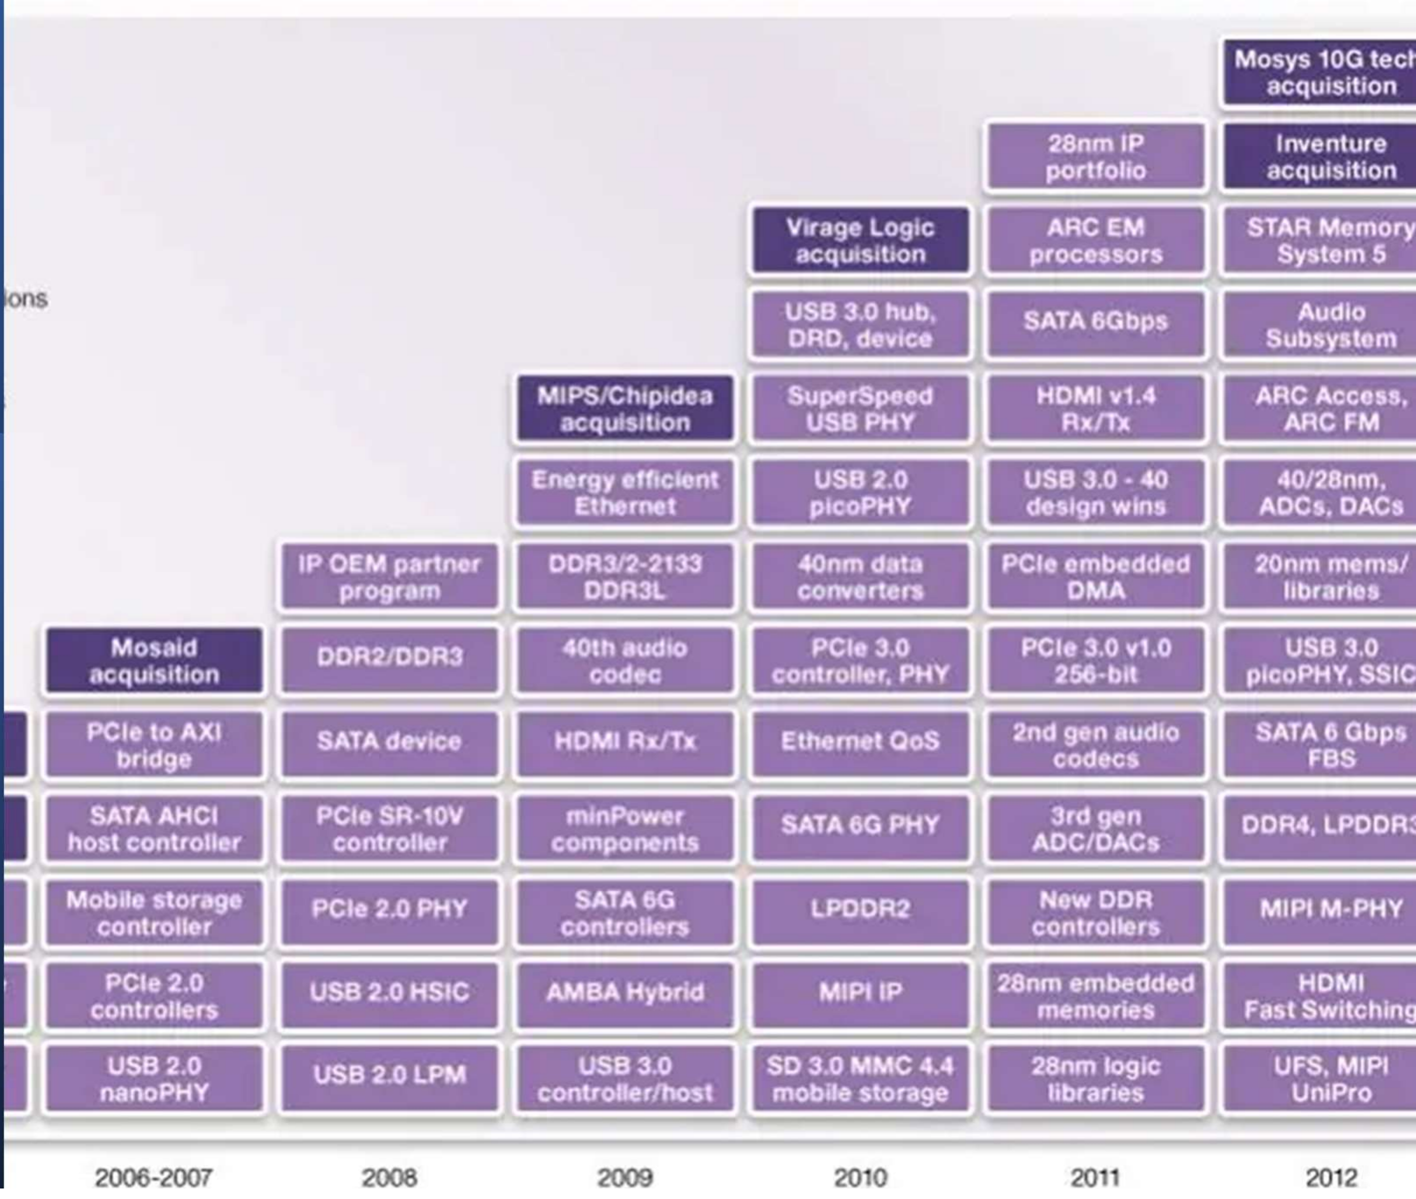

ions

					Mosys 10G tech acquisition
				28nm IP portfolio	Inventure acquisition
			Virage Logic acquisition	ARC EM processors	STAR Memory System 5
			USB 3.0 hub, DRD, device	SATA 6Gbps	Audio Subsystem
		MIPS/Chipidea acquisition	SuperSpeed USB PHY	HDMI v1.4 Rx/Tx	ARC Access, ARC FM
		Energy efficient Ethernet	USB 2.0 picoPHY	USB 3.0 - 40 design wins	40/28nm, ADCs, DACs
	IP OEM partner program	DDR3/2-2133 DDR3L	40nm data converters	PCIe embedded DMA	20nm mems/ libraries
Mosaid acquisition	DDR2/DDR3	40th audio codec	PCIe 3.0 controller, PHY	PCIe 3.0 v1.0 256-bit	USB 3.0 picoPHY, SSIC
PCIe to AXI bridge	SATA device	HDMI Rx/Tx	Ethernet QoS	2nd gen audio codecs	SATA 6 Gbps FBS
SATA AHCI host controller	PCIe SR-10V controller	minPower components	SATA 6G PHY	3rd gen ADC/DACs	DDR4, LPDDR
Mobile storage controller	PCIe 2.0 PHY	SATA 6G controllers	LPDDR2	New DDR controllers	MIPI M-PHY
PCIe 2.0 controllers	USB 2.0 HSIC	AMBA Hybrid	MIPI IP	28nm embedded memories	HDMI Fast Switching
USB 2.0 nanoPHY	USB 2.0 LPM	USB 3.0 controller/host	SD 3.0 MMC 4.4 mobile storage	28nm logic libraries	UFS, MIPI UniPro
2006-2007	2008	2009	2010	2011	2012

Networks on Chip

Systems Composition

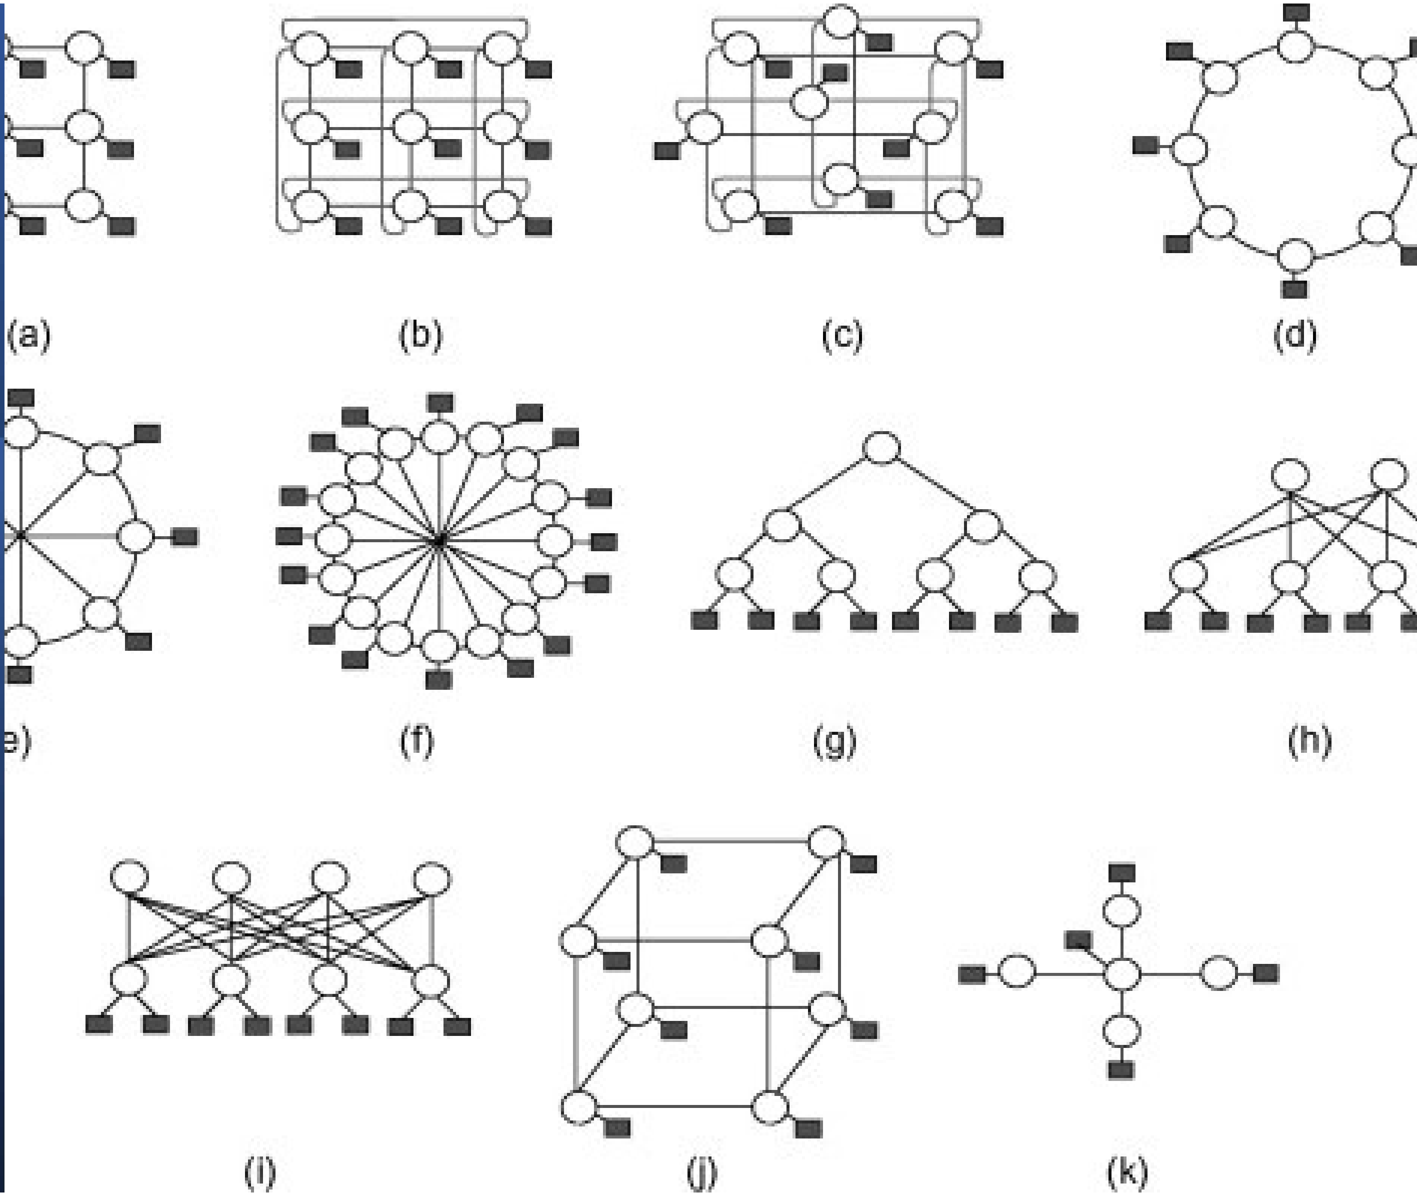

https://www.sciencedirect.com/science/article/abs/pii/S0141933109000374

Systems on Chip SOC

Almost everything on one Die!

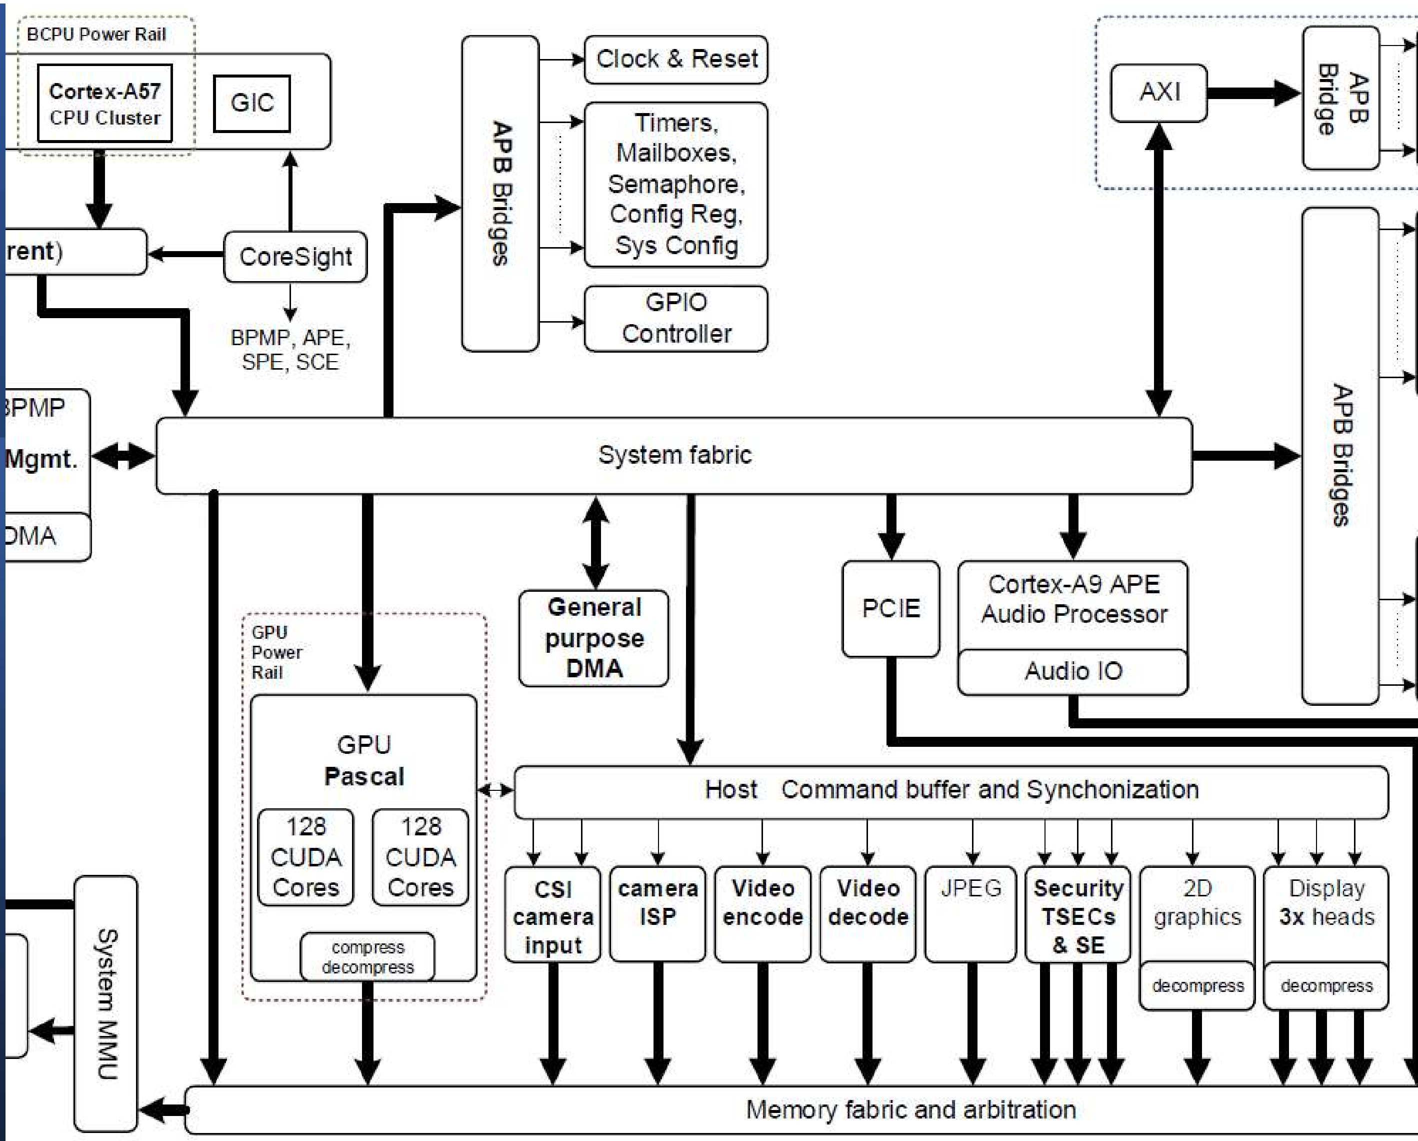

https://developer.nvidia.com/blog/jetson-tx2-delivers-twice-intelligence-edge/

Internet Users Worldwide

1995 *WWW Introduced*

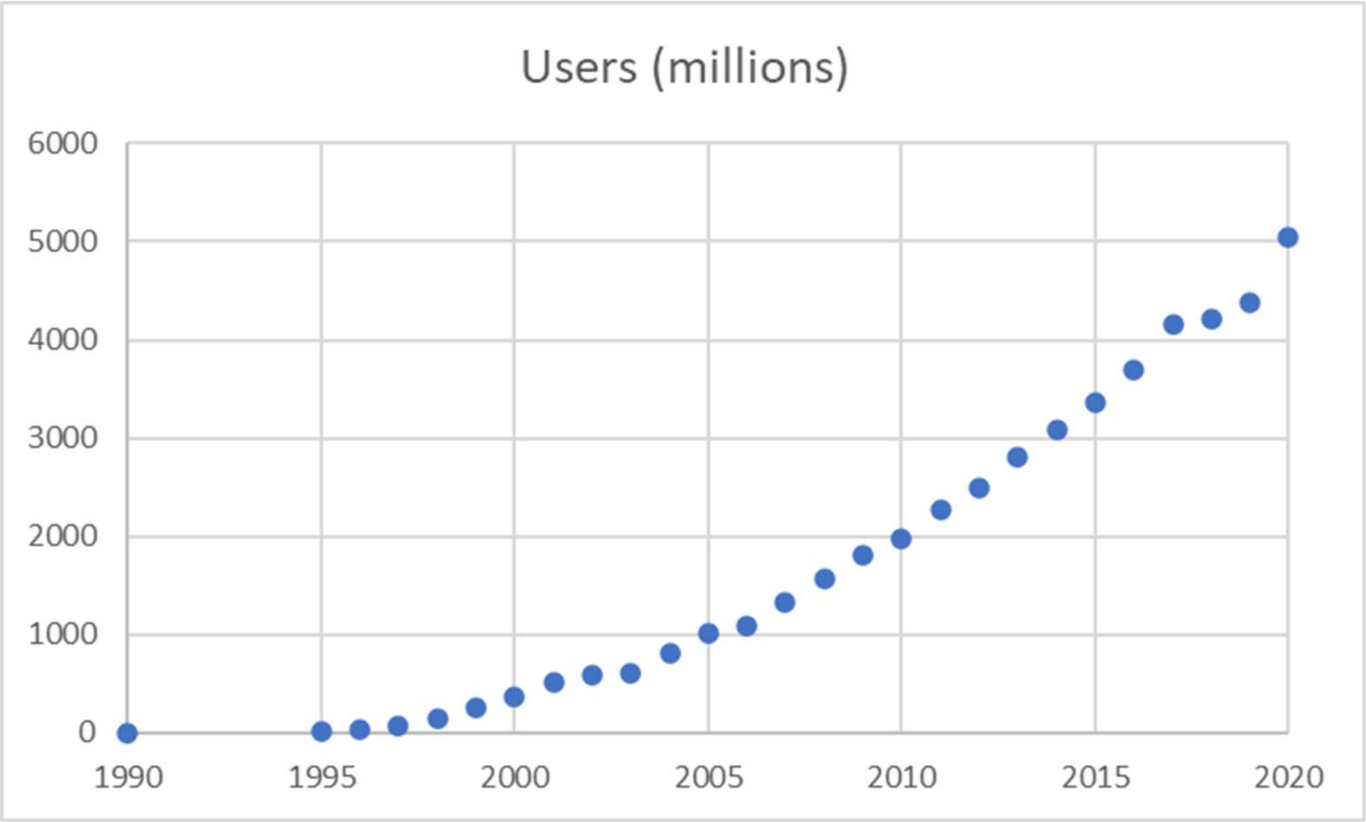

Data Source: https://www.internetworldstats.com/emarketing.htm

Big Routers for Internet Expansion

Cisco GSR 12000

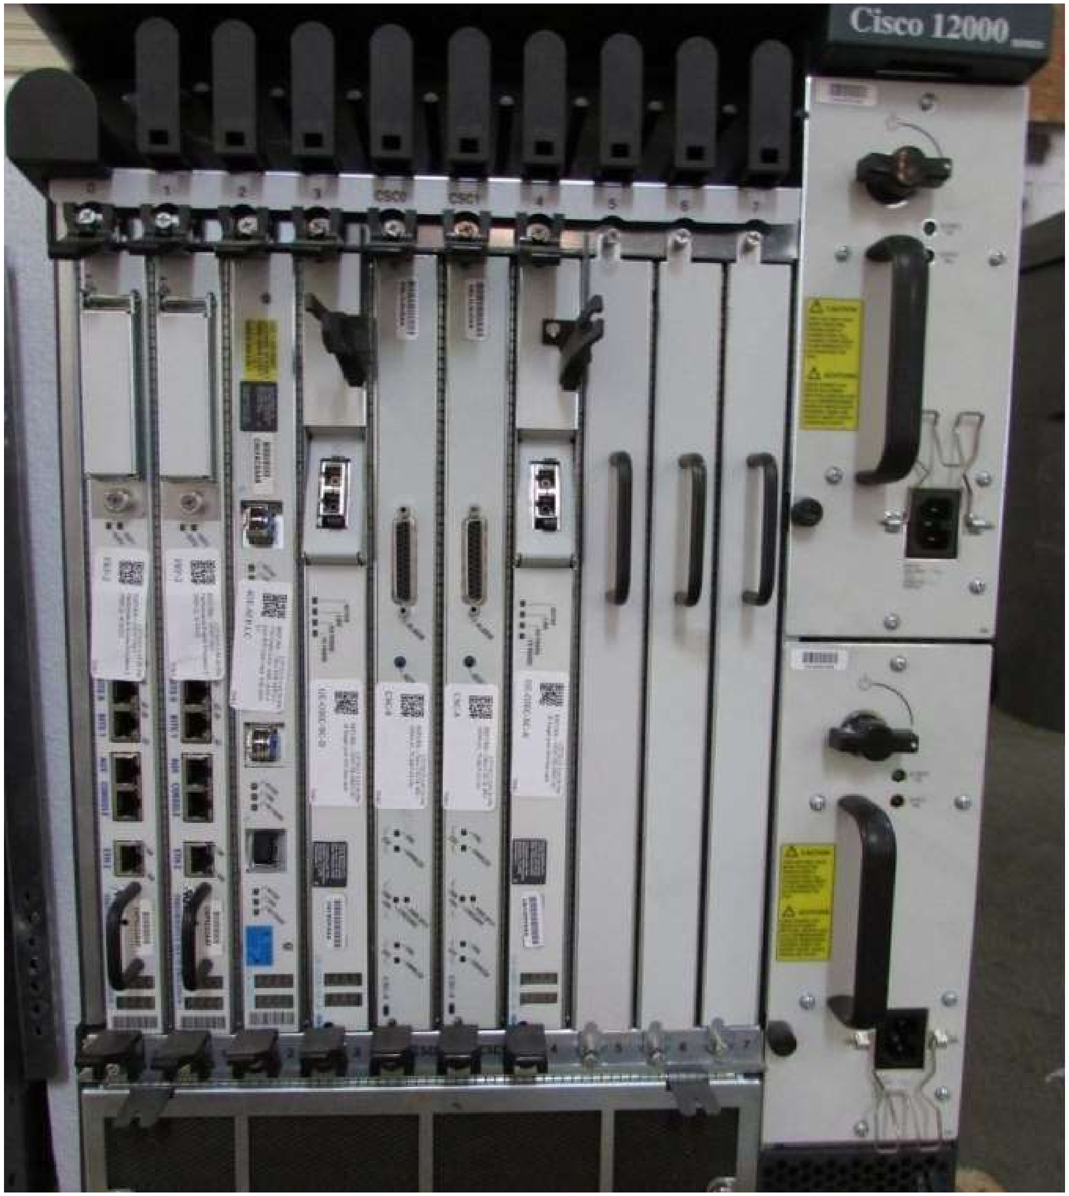

Field Programmable Gate Arrays FPGA

Another dimension of flexibility

Reconfigurable Computers in Data Centers

Application Acceleration On-Demand

SoC Emulator in the Cloud

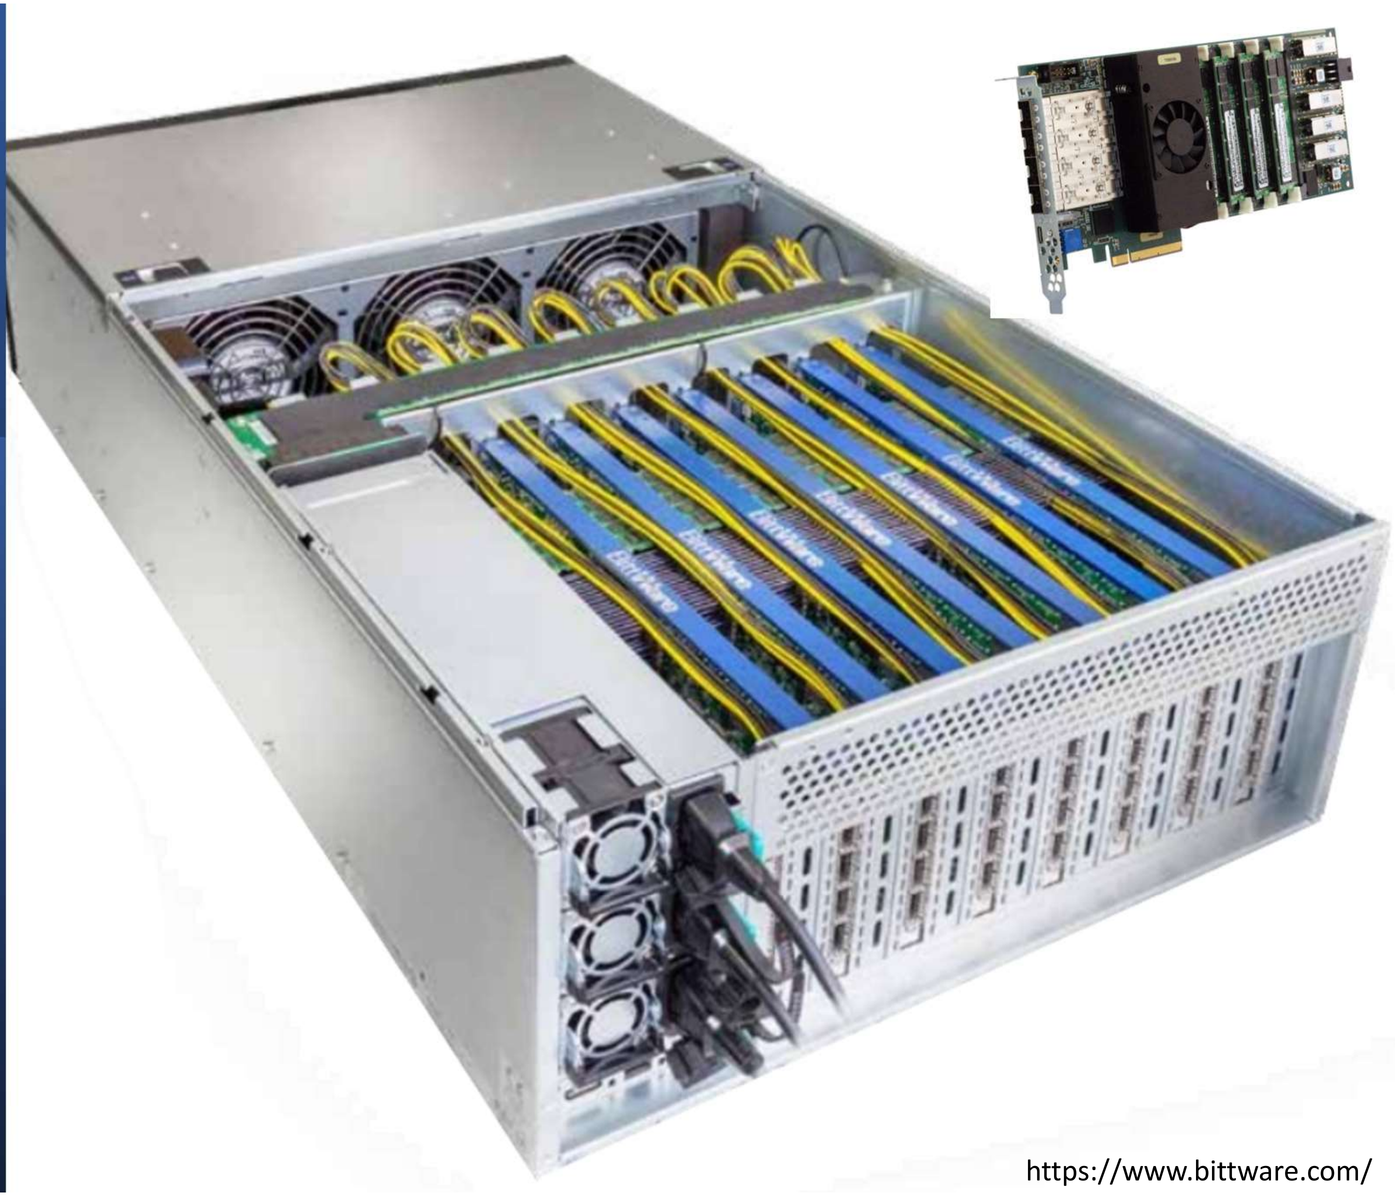

https://www.bittware.com/

Hyperscale

@Google

Gen4 TPU Supercomputer Pods

Deep Learning at Scale

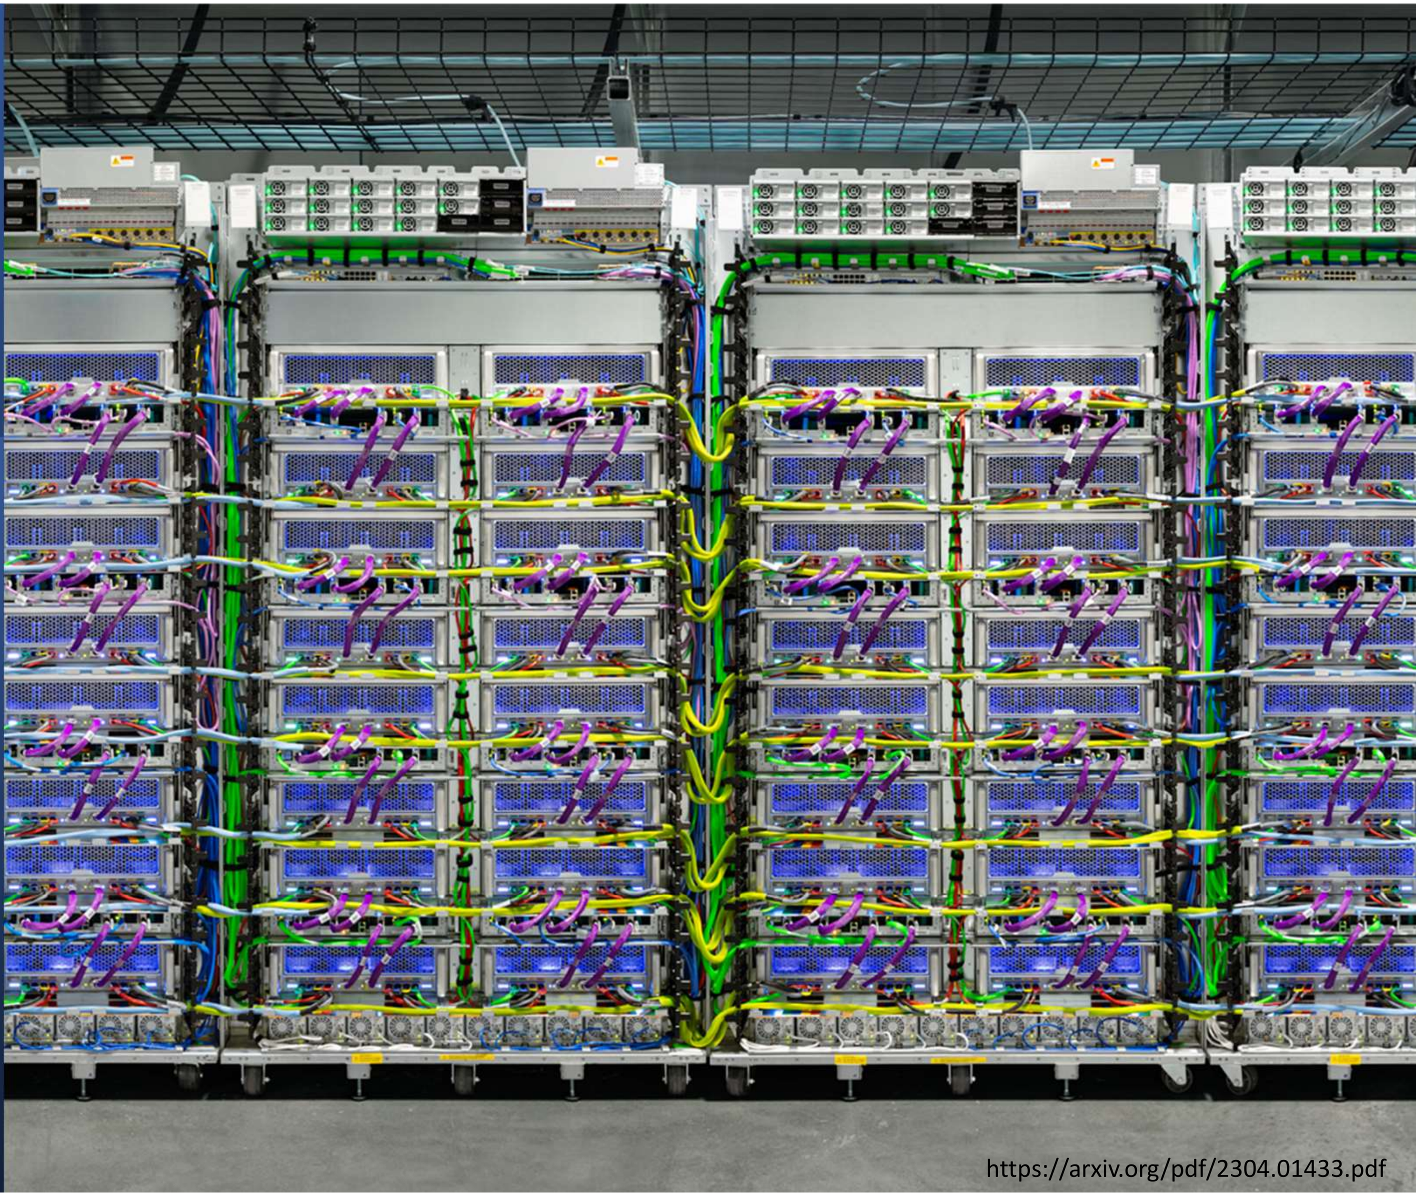

https://arxiv.org/pdf/2304.01433.pdf

Cray-1

Fastest Computer 1976 – 1982

“The world’s most expensive love- seat!”

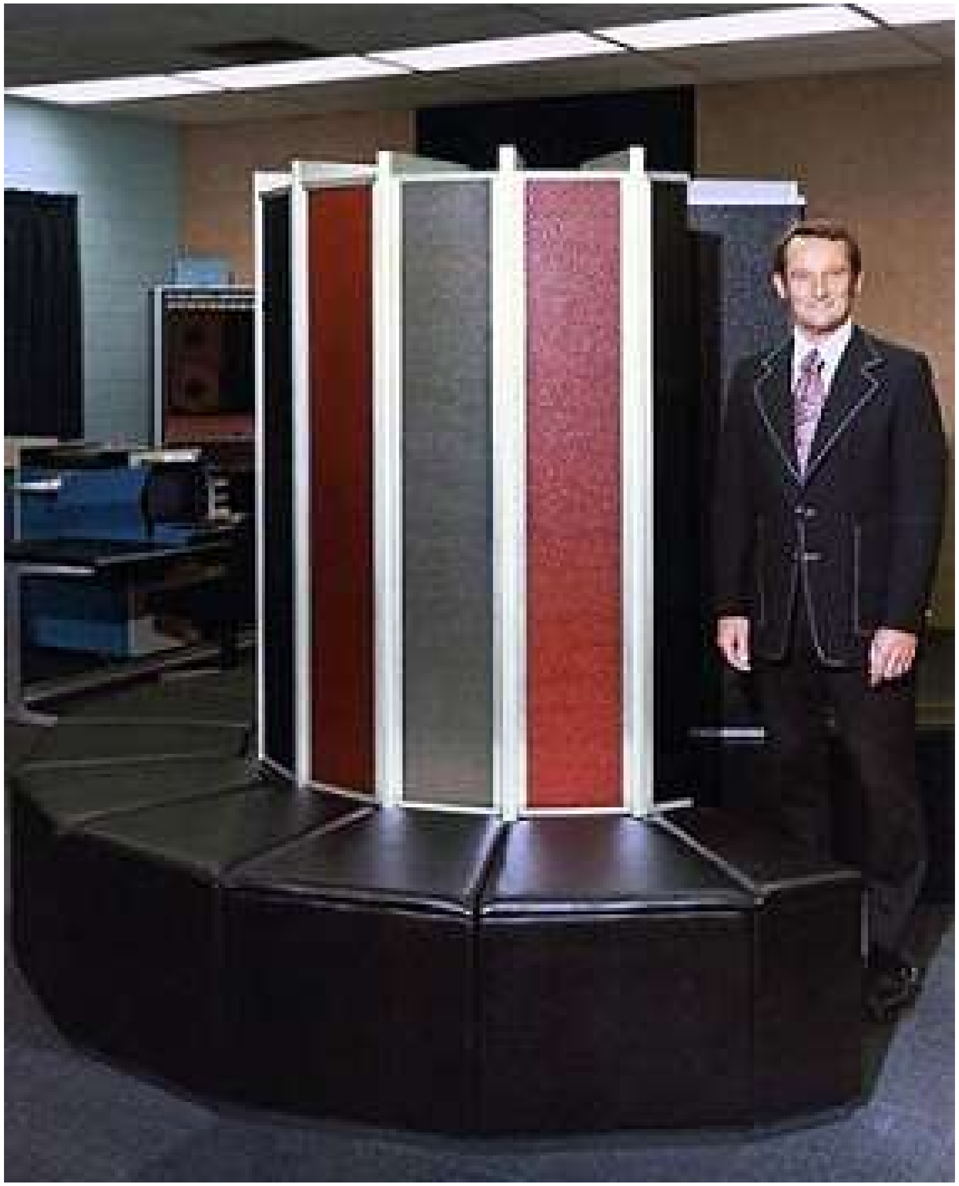

Cray-1 Architecture

SIMD Vector Processor

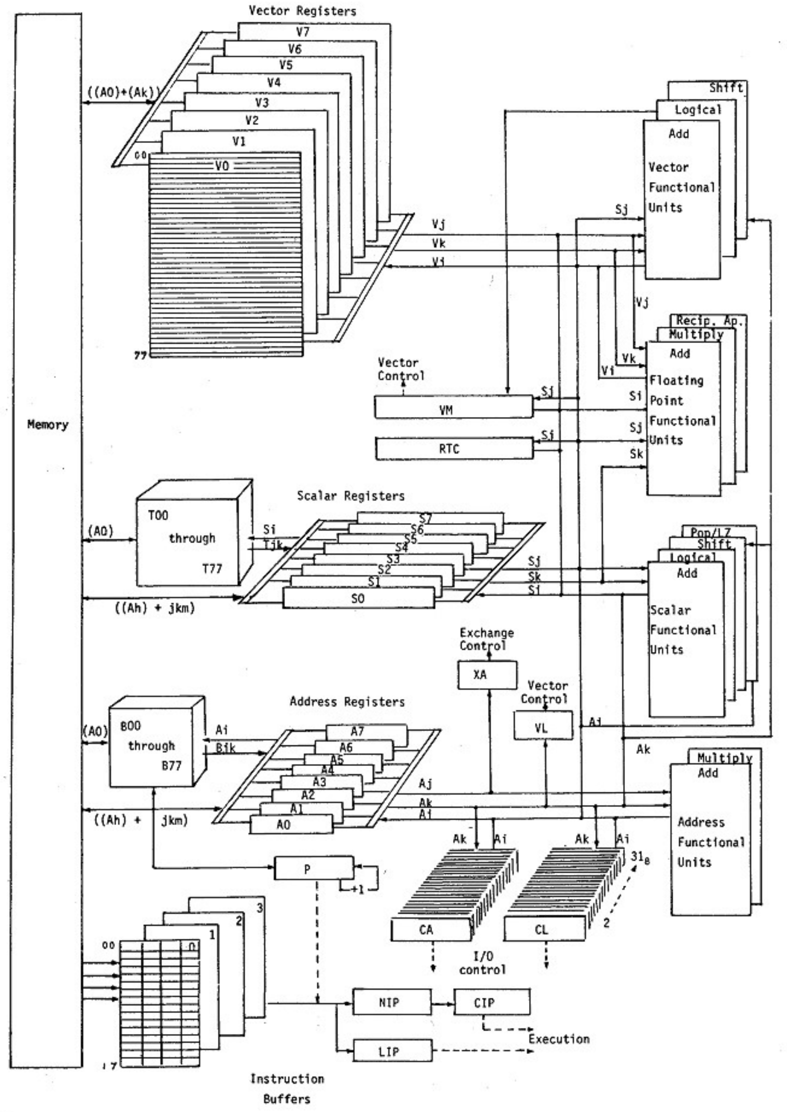

Figure 3-1. Computation section

Gen4 Tensor Processor Unit

Chiplets
Water Cooled
Vector Processing
Networked

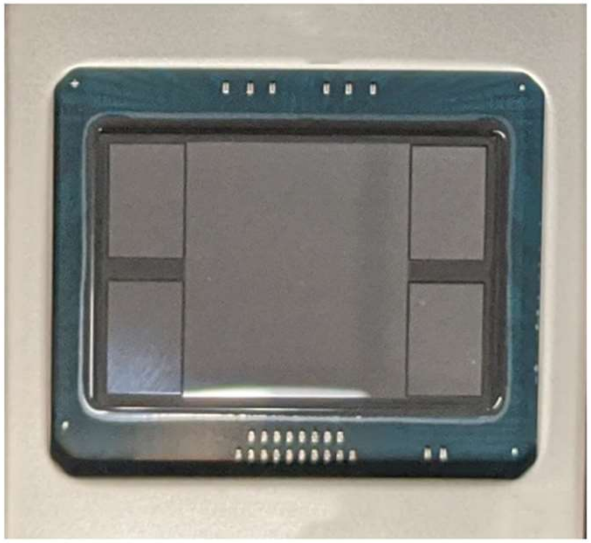

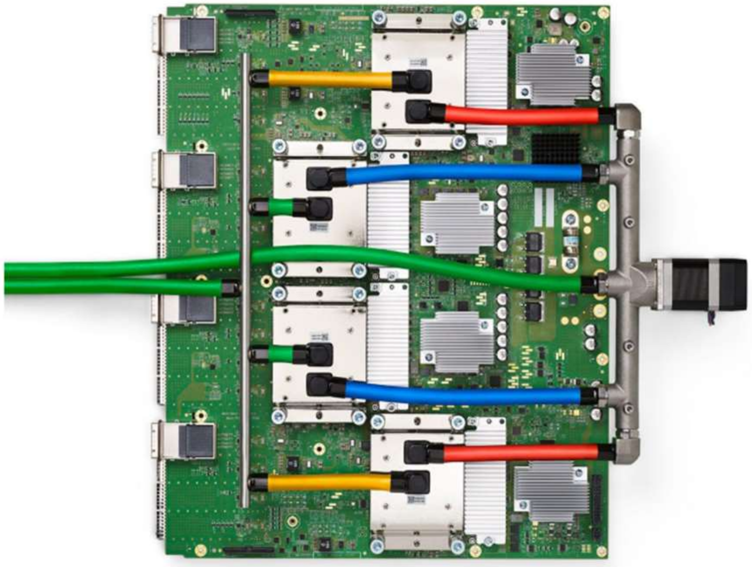

https://arxiv.org/pdf/2304.01433.pdf

Density & Power Optimization with Multi Chip Packaging

Exploit Old Idea

Chiplets on Silicon Interposers

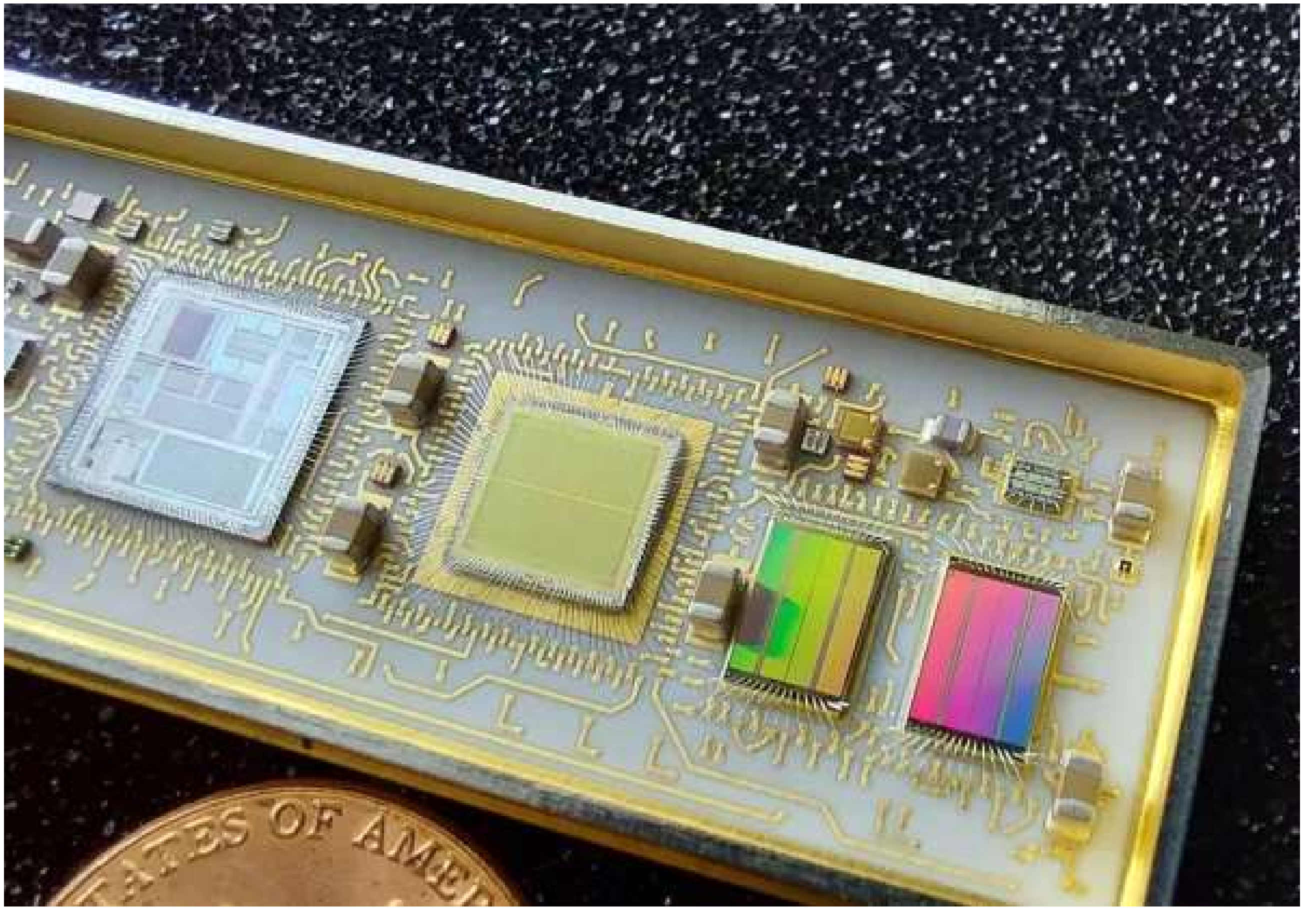

https://www.texasmicroelectronics.com/category/product-information/

Interposer-Chiplet Stack Up

Intel Ponte Vecchio

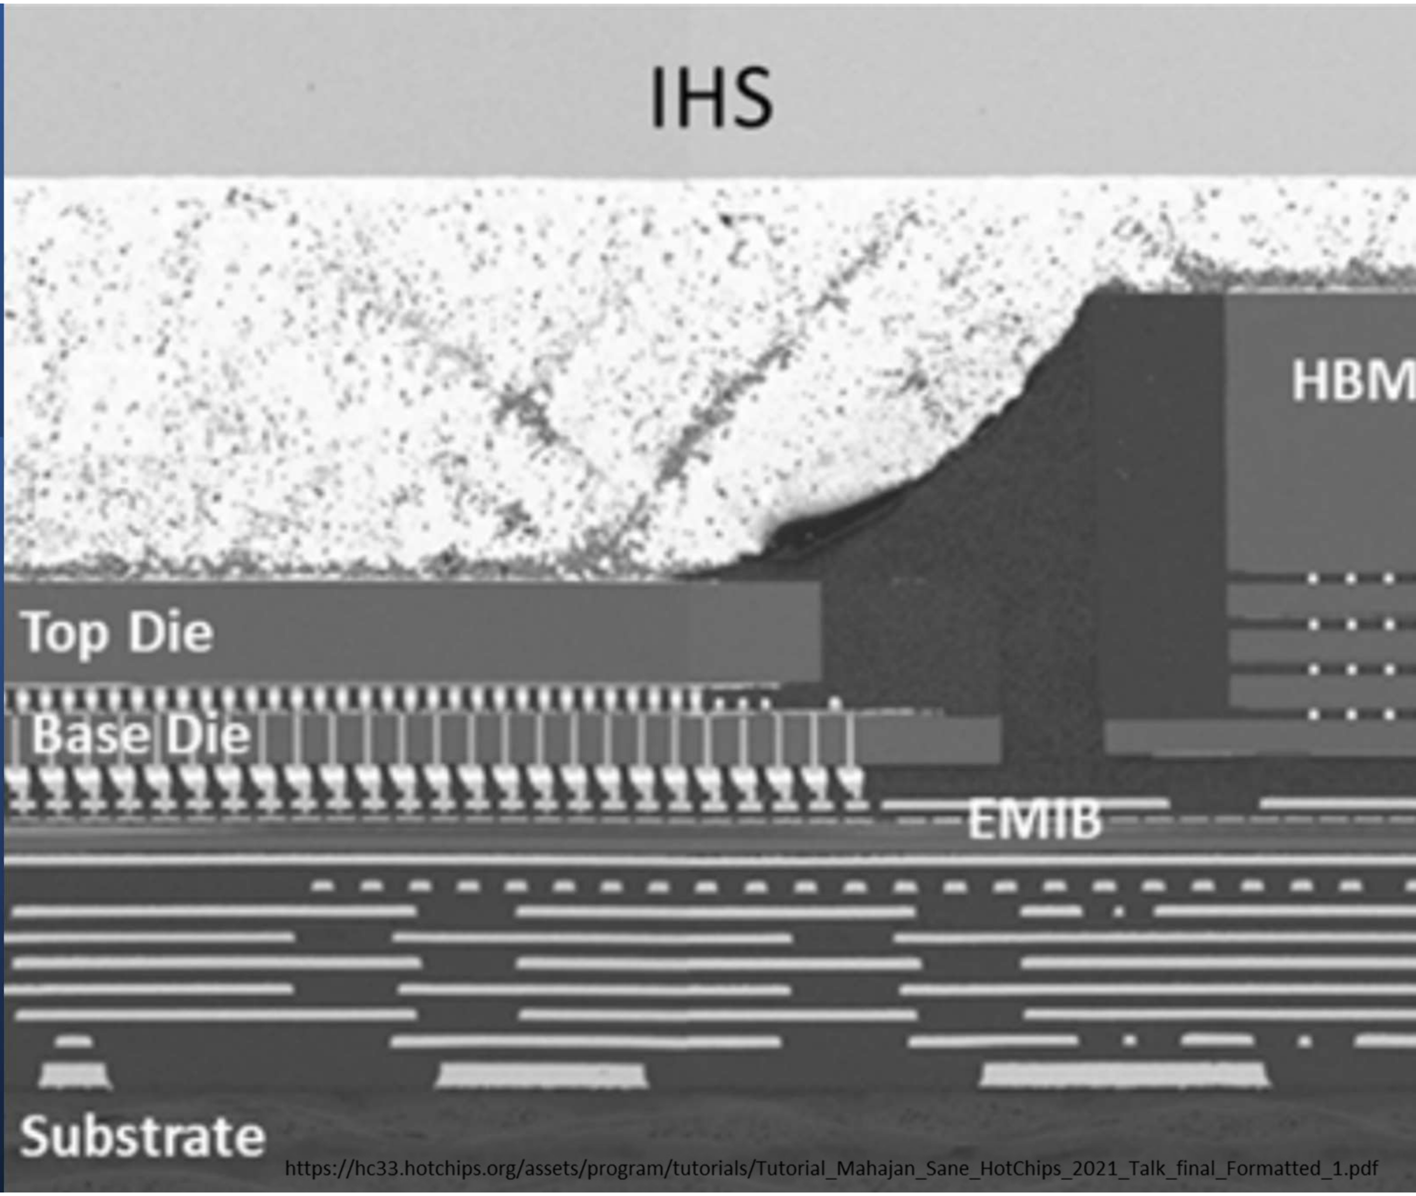

https://hc33.hotchips.org/assets/program/tutorials/Tutorial_Mahajan_Sane_HotChips_2021_Talk_final_Formatted_1.pdf

Protected by Titan

Google Root-of-Trust for the Data Center

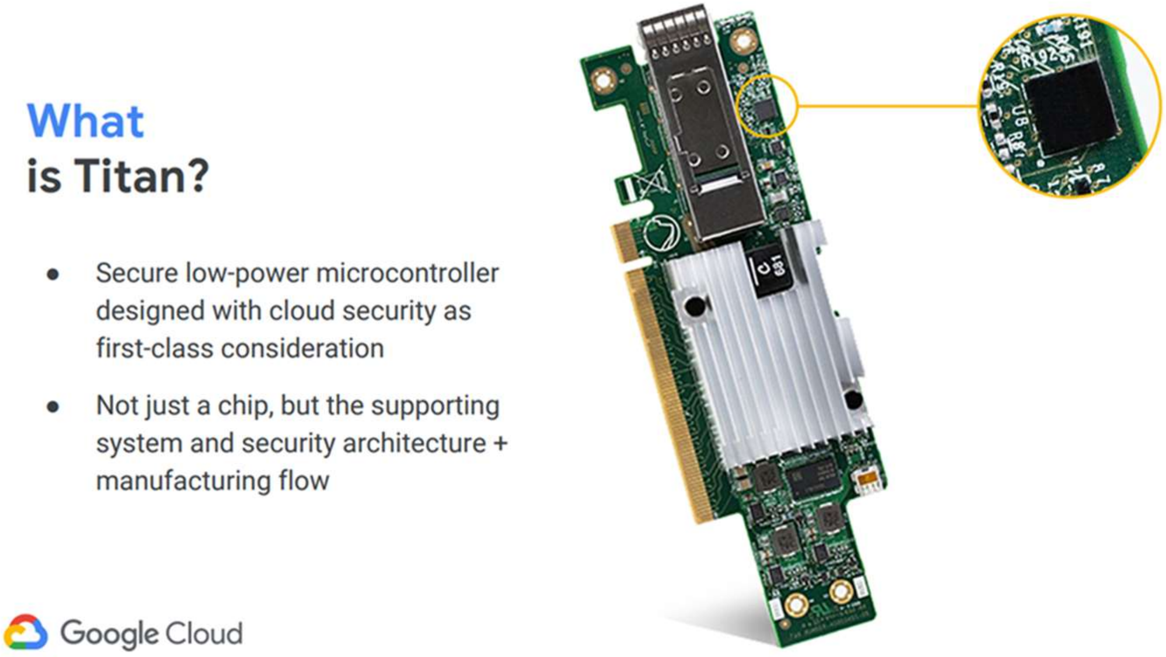

What is Titan?

- Secure low-power microcontroller designed with cloud security as first-class consideration
- Not just a chip, but the supporting system and security architecture + manufacturing flow

Google Cloud

Titan-M1

Root-of-Trust

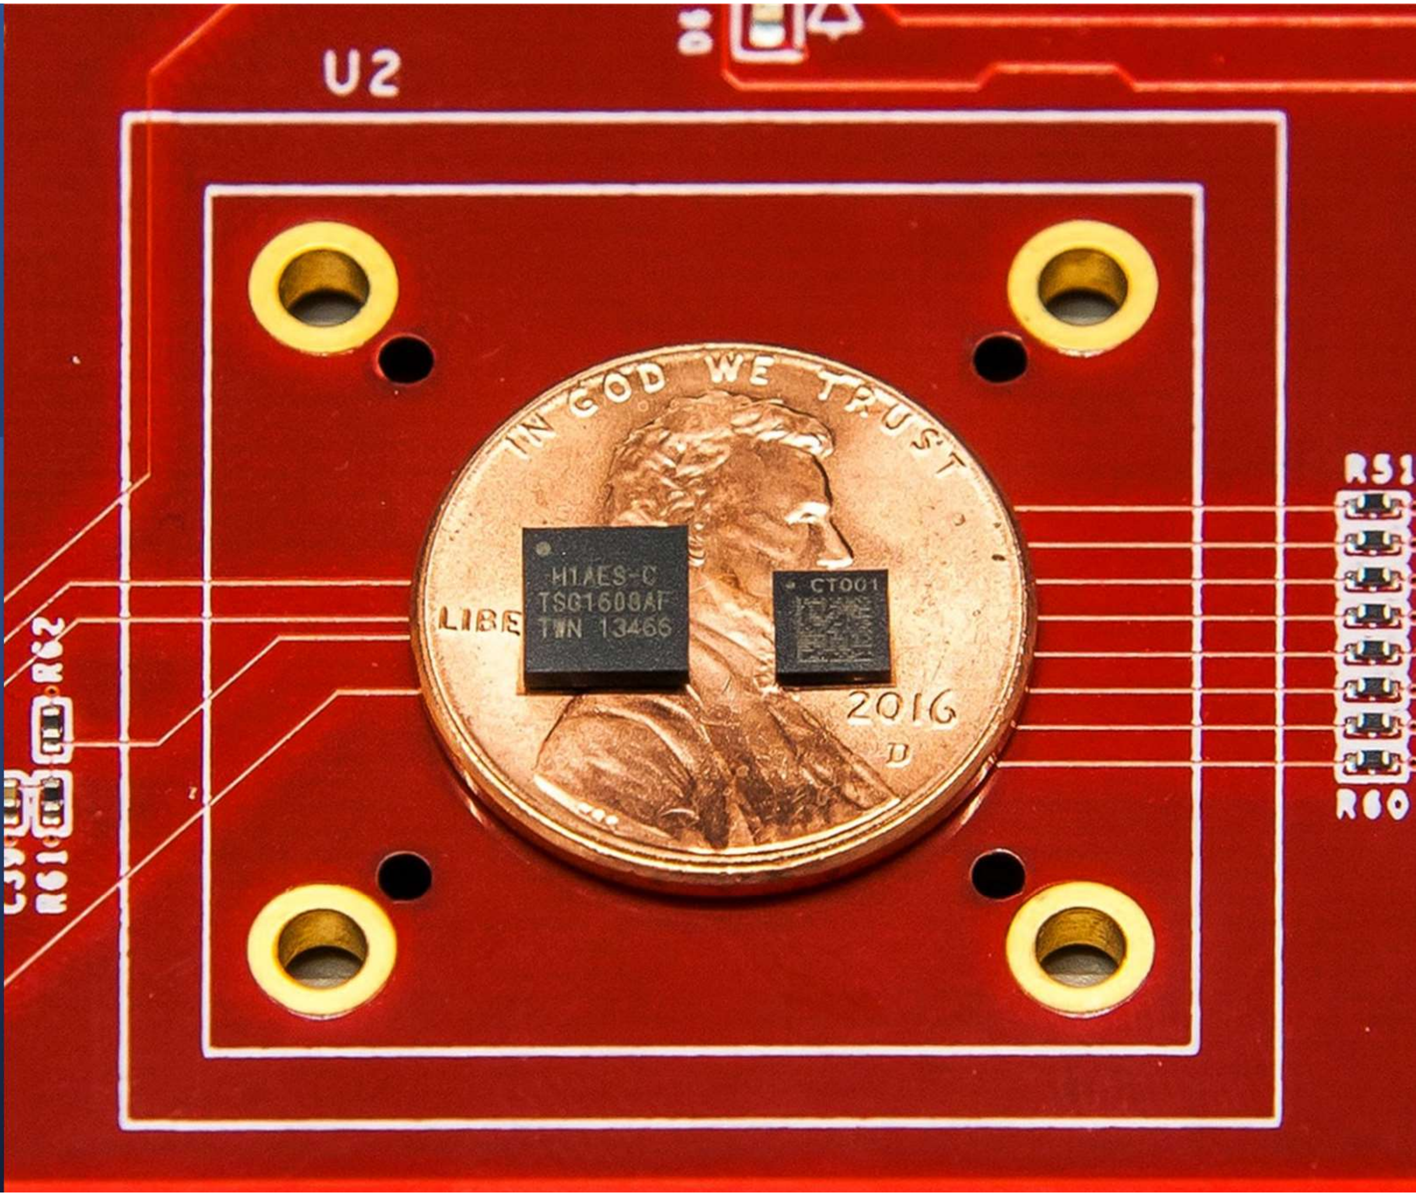

Titan Everywhere

Everyone deserves good Security!

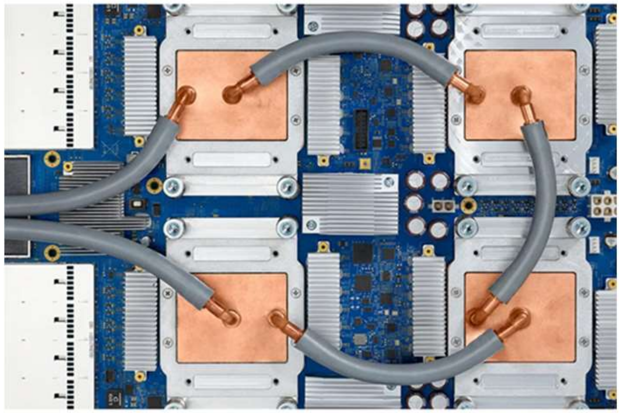

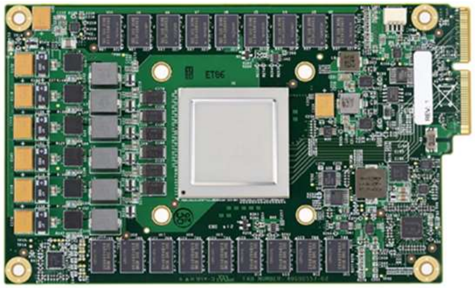

But…
Is it Safe?

Trust but Verify!

https://cinema-fanatic.com/2011/05/12/movie-quote-of-the-day-marathon-man-1976-dir-john-schlesinger

Side Channel Analysis

Information leakage via Power and EM Signals

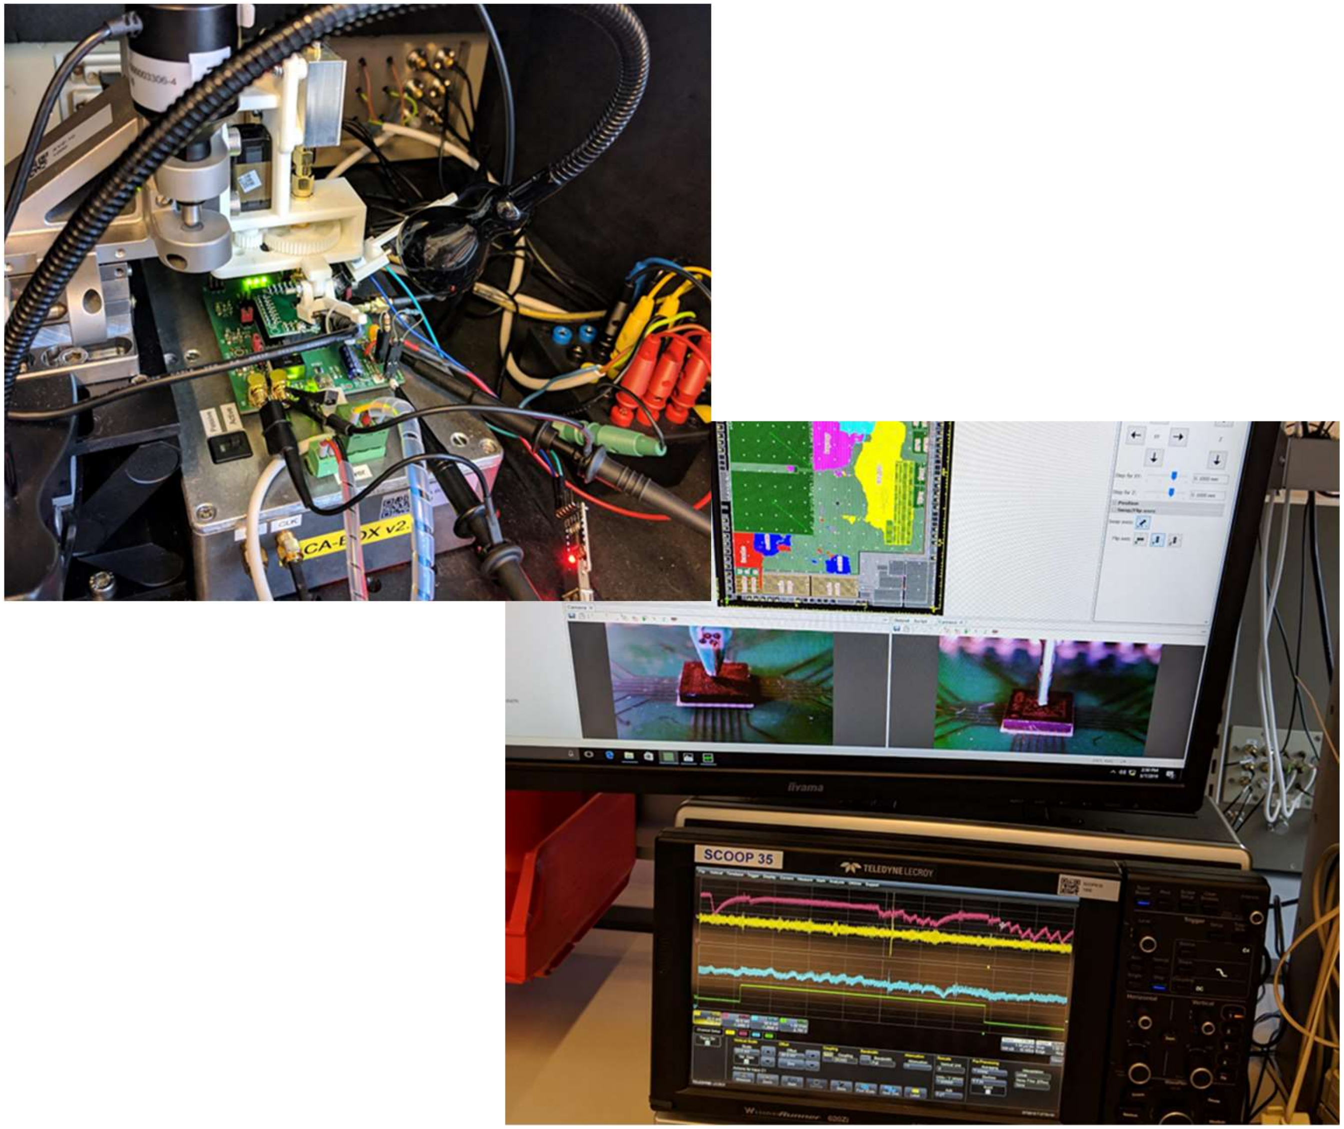

NSA Tempest

1950s onwards

Listening in on "secure" comms 100s of meters distant

https://www.nsa.gov/portals/75/documents/news-features/declassified-documents/cryptologic-spectrum/tempest.pdf

IR Emissions Side Channel

Investigating Chip Floorplan

16nm node
1300nm IR
Resolution
3 micron/pixel

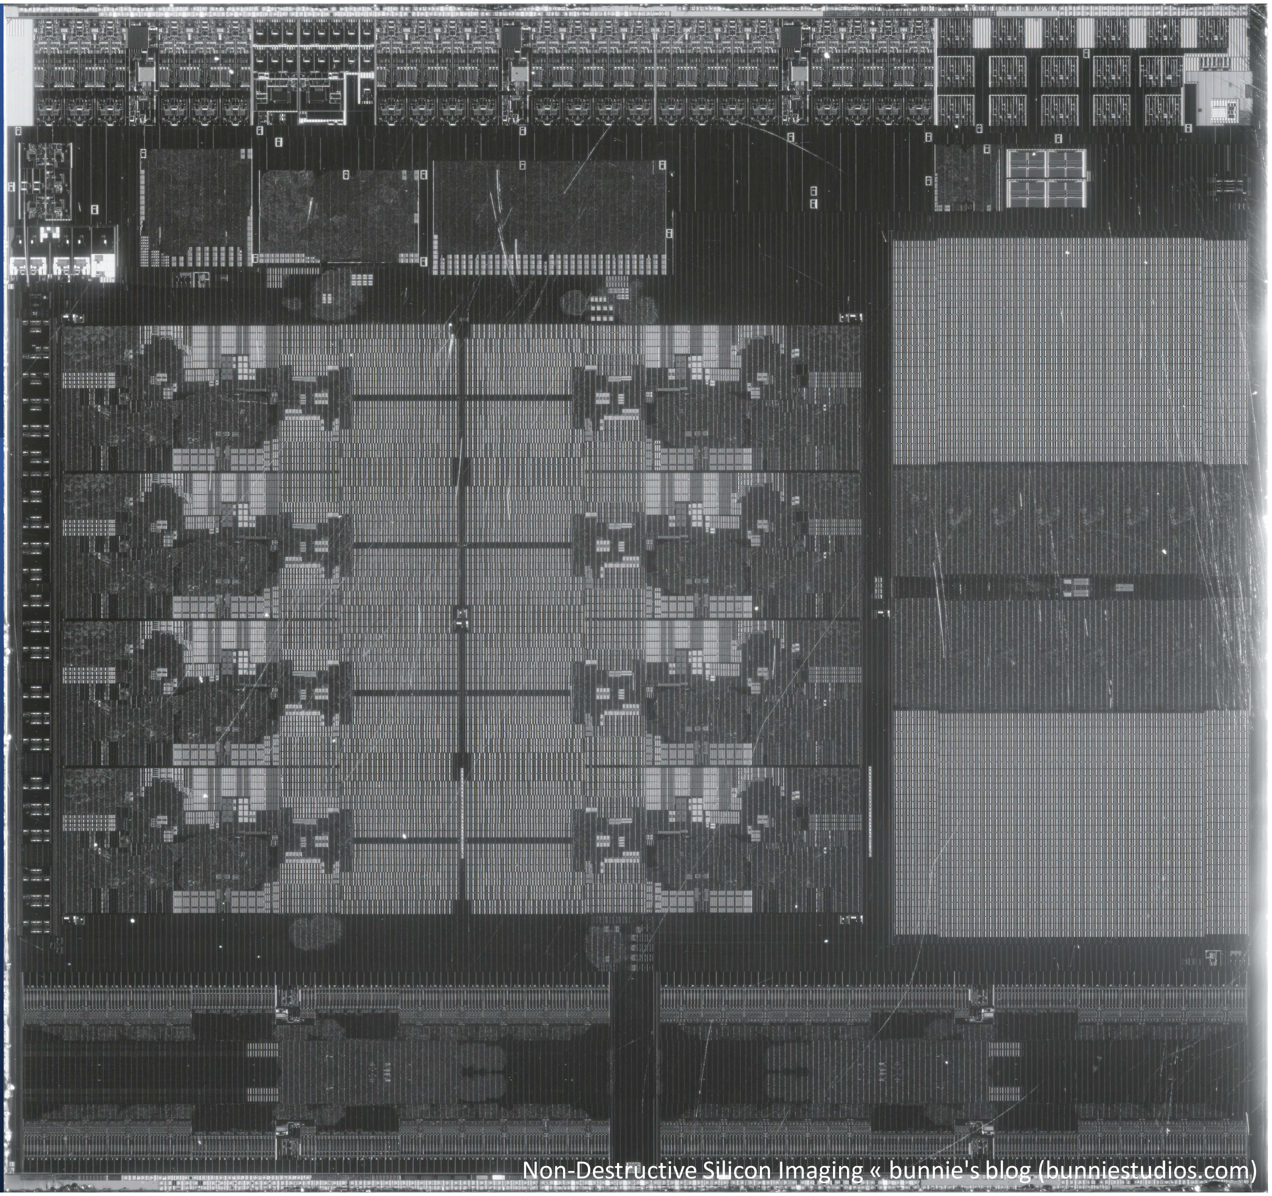

Non-Destructive Silicon Imaging « bunnie's blog (bunniestudios.com)

IR Emissions

Read the Data

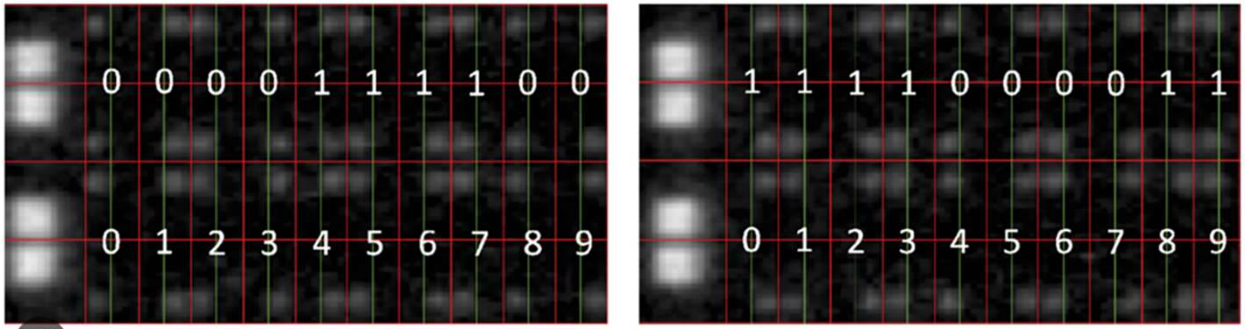

(a) Checkerboard design.

(b) Inverted checkerboard design.

https://www.semanticscholar.org/paper/Direct-read-of-idle-block-RAM-from-FPGAs-utilizing-Couch-Whewell

Some Conclusions …

- History repeats
  - But with new variations
- Abstraction
  - Helps us deal with complexity
  - But abstraction also obscures
- Lower levels can provide valuable insights
  - Should not be considered “fixed” forever
- Explore and Recycle
  - Changing constraints yields different results

Thank You!

HARRIS 2024 Keynote Abstract

Perspectives from Four Decades of Chip Design

The past four decades have seen a dramatic evolution of chip design technology. We've gone from 5 micrometer NMOS down to 3 nanometer CMOS with a corresponding multi-millionfold growth in transistor density. The regular introduction of new abstraction layers and hardware microarchitectures supported by EDA design tools has enabled the industry to deal with ever increasing complexities. This trend does not appear to be ending anytime soon. We take a retrospective view of how we got to the current state of the art and find there are recurring patterns we can use to guide us forward. New opportunities arise to reuse past patterns in novel new ways as the technology constraints shift. Understanding the foundations is key to building the next generation technologies.

Backup

A bit about me …

University of Toronto	***Electrical Engineering***
Santa Clara University	***Applied Mathematics***

Siemens	*EEPROM*
Bell-Northern Research	*DSP*
Synopsys	*Arithmetic IP*
C-Cube	*MPEG Video*
Cisco	*Routers*
Altera	*FPGA*
Google	*Titan Security*
Synopsys	*Crypto IP*

Silicon RE Capabilities

Professional

source: TechInsights

Silicon RE Capabilities

Independent

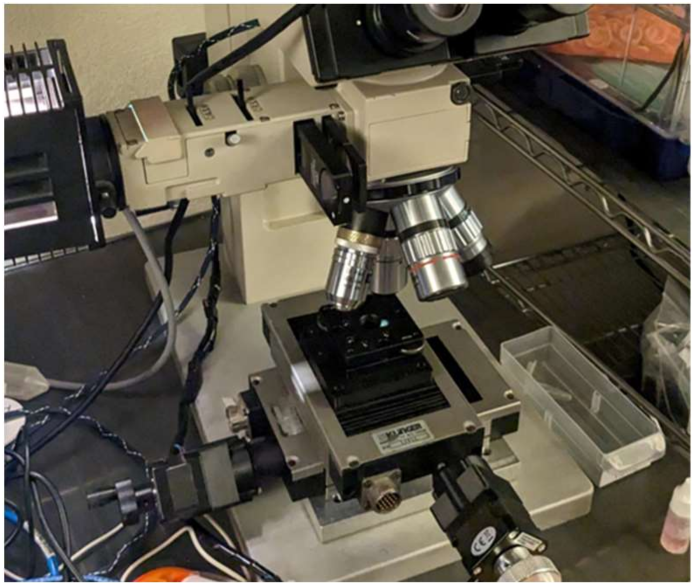

Source: Paul Scheidt @kraftwerklabs.com

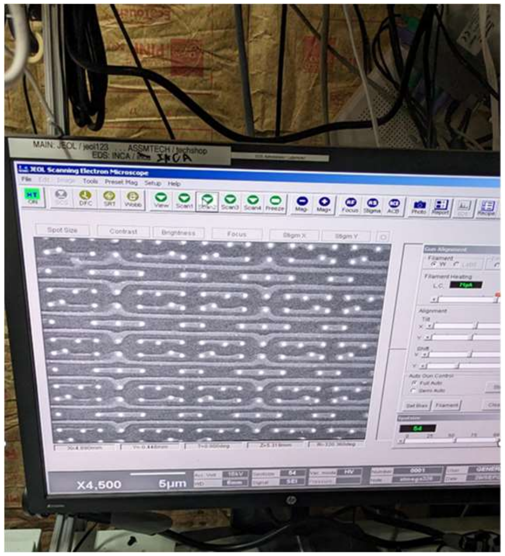

Direct Wafer Probing

Testing & Debugging

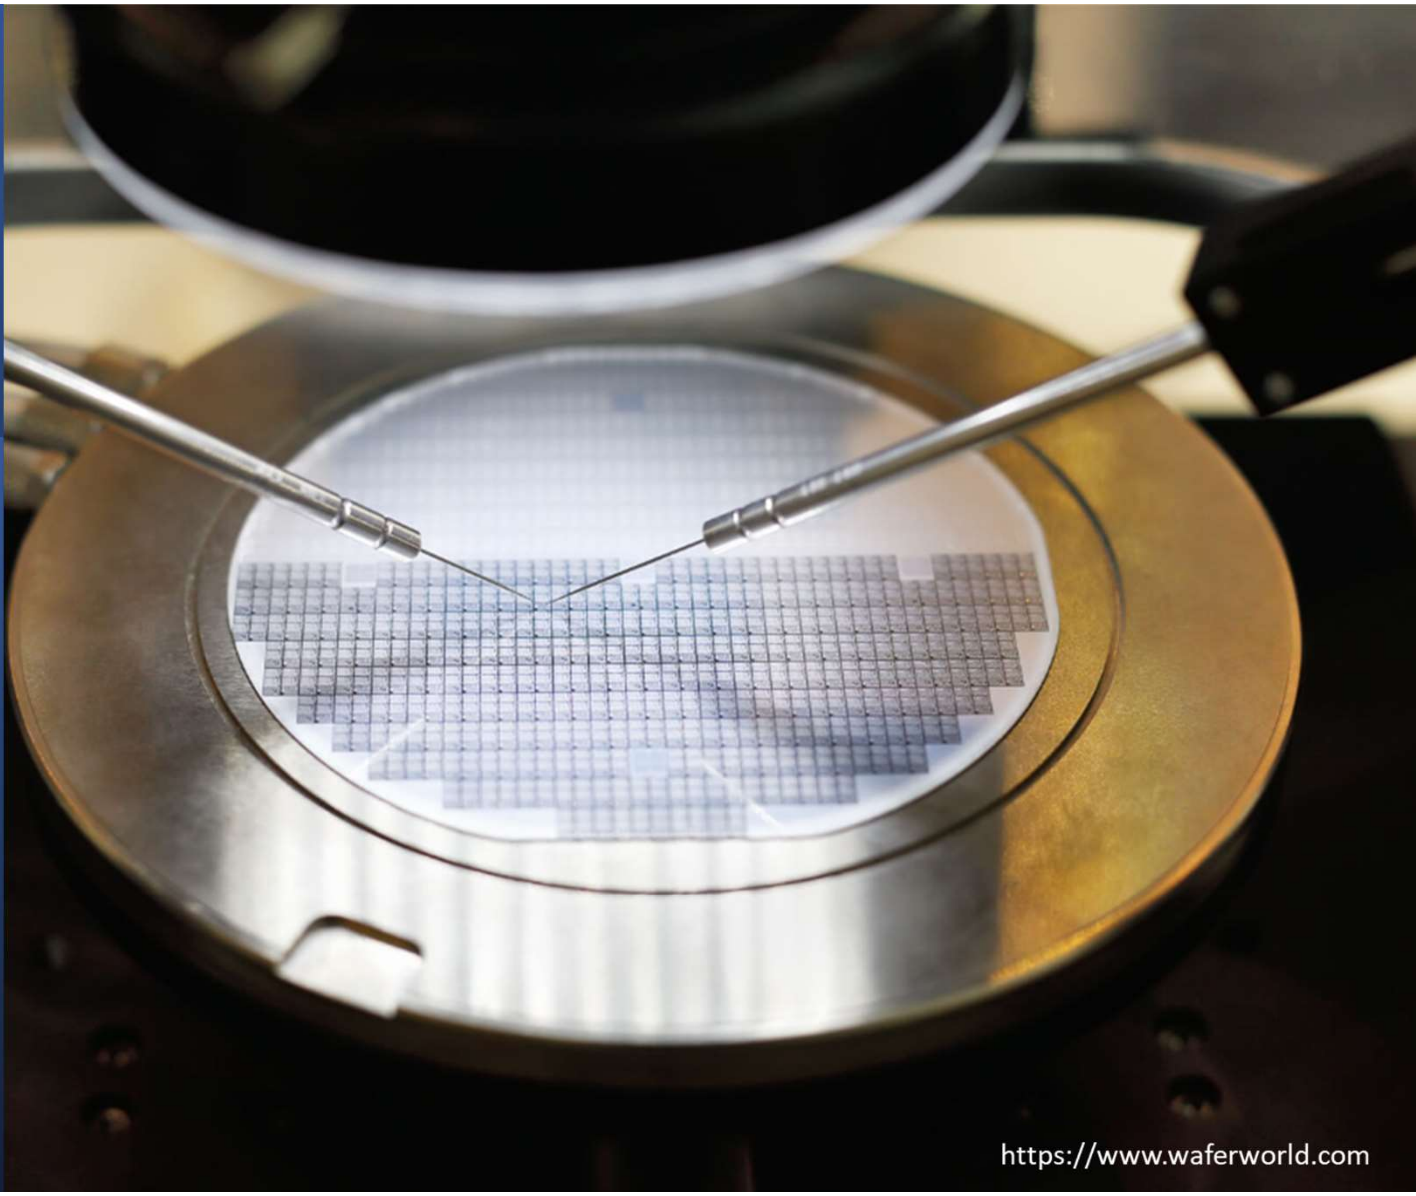

https://www.waferworld.com

DIY Probing

What's happening inside there?

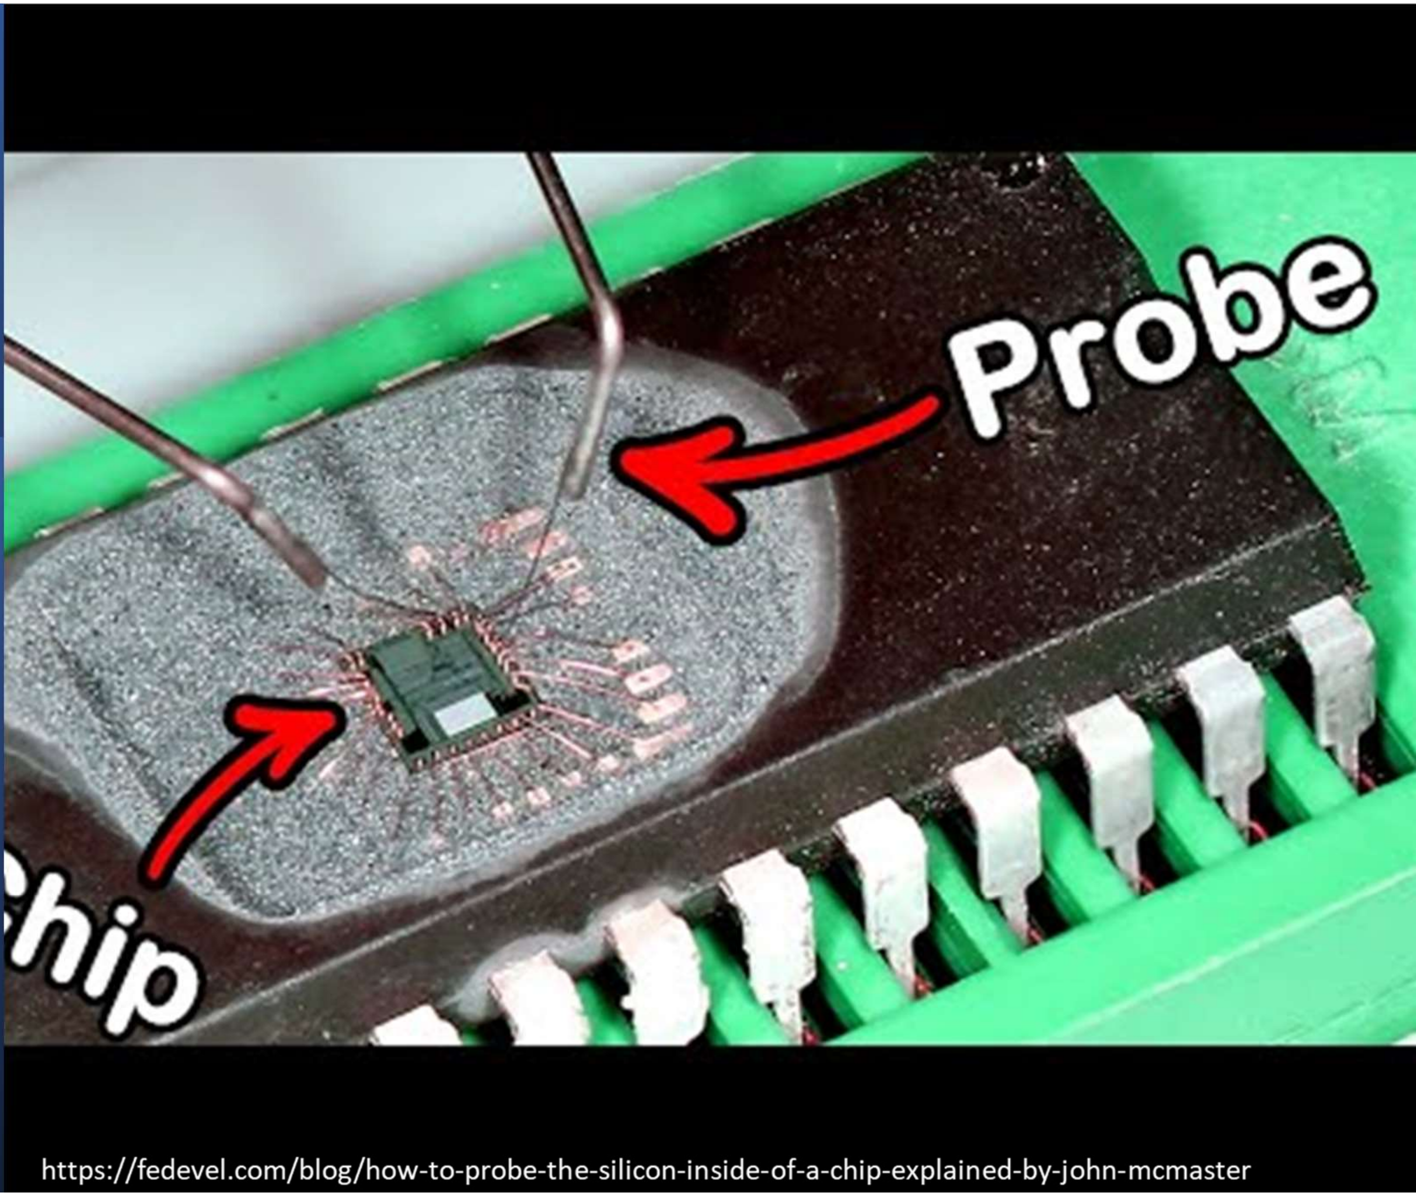

https://fedevel.com/blog/how-to-probe-the-silicon-inside-of-a-chip-explained-by-john-mcmaster

Logic Circuit Design Evolution

Multiplexors

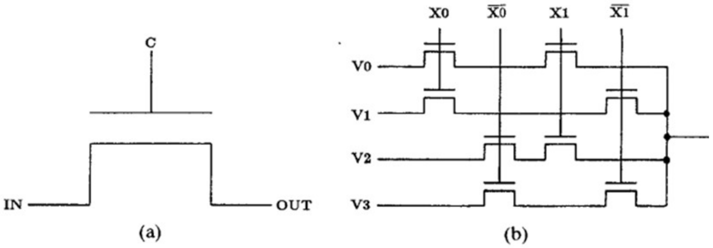

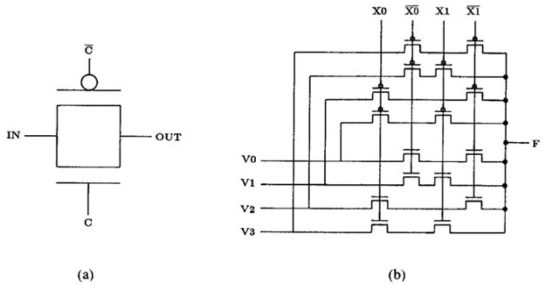

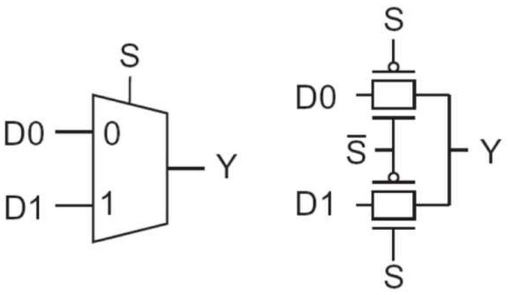

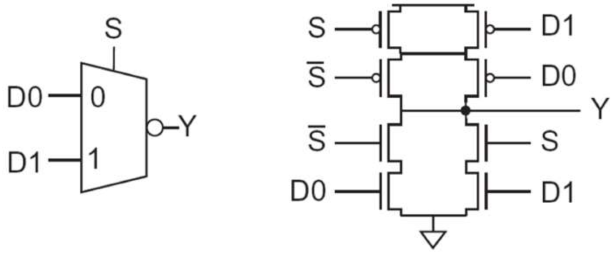

IEEE Journal of Solid State Circuits Vol SC-22, No. 2, 1987
Cornell ECE5745 CMOS Circuits

RE

Threats & Opportunities

The Good

- Security Audit
- Increased Trust

The Bad

- Intellectual Property Theft
- Cloning

The Ugly

- Supply Chain Attack
- Trojans

- T