Lecture-1 (Logistics and Introduction) CS422-Spring 2018

Biswa@CSE-IITK

Instructor

Biswa ~~(Biswabandan, Sir, Prof., Dr., Er., *-Biswa)~~

Sir/Prof./…… outlawed with CS422 and Biswa

Website: http://www.cse.iitk.ac.in/biswap

Contact: KD 203, biswap@cse.iitk.ac.in
Office Hours: Friday, 12 noon

Teaching and Research Interests:
Computer Architecture, Arch-OS interface, and Systems Security

Course Staff: BASS

Biswa

Ayushi

Computer Security

Salman

Memory Side-channel attacks

Shubham

Data Mining, Computer Architecture, Cyber Security

Logistics

When: Mon/Wed. 14.00-15.30 Hrs,
Where: KD 101, What: You know it

Course website: www.cse.iitk.ac.in/~biswap/CS422.html

Piazza: For online discussions

Submission of assignments: Canvas

Register/Drop ASAP (if interested/not interested)

CS422: Two states ?

Curious mind: Eager to learn, understand, implement, and analyze

Non-curious mind: Not eager to learn, understand, implement, and analyze. Interested in writing an exam and getting a grade for CS422.

Which state are you in?

Assessment Policies – Curious Minds

20 = 10 (in-class online tests) * 2

10 = Mid-term

20 = End-term

05 = Paper presentation

05 = Research paper review

40 = 4 programming assignments * 10
(-10%) = Late submission days
(+10%) = Early submission (at least 24 hrs before)

-1 : coming late to the lectures OR
asking for deadline extensions for any assignments

-1 : referring Biswa as Prof., Sir, Dr., *Biswa*

+1 : Piazza/In-class participation

Assessment Policies: Non-curious Minds

50 = Mid-term

50 = End-term

-1 : coming late to the lectures OR
asking for deadline extensions for any assignments

-1 : referring Biswa as Prof., Sir, Dr., *Biswa*

+1 : Piazza/In-class participation

Group Based Points

-1 : coming late to the lectures OR
asking for deadline extensions for any assignments

-1 : referring Biswa as Prof., Sir, Dr., *Biswa*

+1 : Piazza/In-class participation

Group C

Group NC

Why Two Policies ?

I do not want you to become exam-takers. Instead, I want you to learn.

"I never teach my pupils, I only provide conditions in which they can learn" – Who ? Guesses

"I will not teach CS422, I will only provide different ways in which you can learn CS422" – Who ? Of course Biswa

So be curious, eager to learn (not through rote learning).

CS 422 Assessment

20 = 10 (in-class online tests) * 2

Fundamentals

10 = Mid-term

20 = End-term

Application of fundamentals

05 = Research paper review

05 = Paper presentation

Learn what others are doing and teach others what you know

40 = 4 programming assignments * 10

1. Learn to understand
2. Learn to design existing ones
3. Learn to analyze
4. Learn to design new ones

Openness

In class:
2 + 2 = 4

Closed *

Homework:
2 + 4 + 2 = 8

Open but close-ended

Exam:
John has 4 apples, his train is 7 minutes early. Calculate the mass of the sun.

* Open *

What I Expect from You?

No open-screens (no nomophobics): No open smart-phones (phones) & laptops/tablets. Keep your phones in silent mode

Open-screens will affect (distract) you, your friends, and me

Ask questions & participate in in-class discussions (worth bonus points)

Paper reading and writing reviews/reports

Understand, implement, and analyze ideas (Hard work and honesty)

Slides will not contain everything. So attend lectures.

What I Expect from You?

Timing	Classes start at 2 PM, not 2.10/15 PM
Cheating	In any form will lead to zero points. Grade will be capped down (one level). To prevent capping down, you have to build architectural tools.
Dropping CS422	Not allowed after Jan 12th 2018. Drop the course before that. Why? It will affect your group points.

What I Expect from You?

Ditch your excuses.

Participate in class/Piazza regularly.
Do not fear about your doubts. Just communicate.
We (you, T.A., and me) will try our best to address it.

Just shout if you do not like something about me or about the course. However, be on the right side and then shout.

What Curious Minds Can Expect From Me?

I will give my 100% in delivering lectures, clearing your doubts, helping you learning.

A relationship based on trust and respect. I will be lenient in terms of deadlines (others) for genuine cases.

Available for meetings regularly.

Reference letters in the future.

However, there is no offence if you are not curious. Perfectly fine.

Questions ??

CS422 In Four Modules – Four Talks

Module 1: Processor

Module 2: Caches

Module 3: DRAM

Module 4: Advanced Topics

Intel ?
Qualcomm?

Advanced Topics (If Time Permits)

Introduction and logistics	Case study: Recent Intel, ARM, AMD, IBM, and Qualcomm based systems	
ISA Tradeoffs, RISC/CISC, RISC-V + Pin Tool	Emerging Memory Technologies	Talk 3 - Vivek Seshadri, MSR ?
Simple pipelining	Network on Chip	
Complex pipelining, Performance Metrics, Amdahl's Law	Power, Energy, and Dark silicon	
More on pipelining, Interrupts, traps, and exceptions	Storage, Heterogenity, Specialization, and Accelaration	
Branch predictors		
Branch predictors + CBP Framework		
O3 Processor, Tomasulo, and ILP		
Brief overview of SMT, VLIW, GPU	Talk 1 - Andre Seznec, INRIA	
Memory Hierarchy, Caches (VIVT, VIPT, etc.)		
Cache Optimizations		
Caches for Multicores, Hardware Prefetching		
Caches for Multicores, Hardware Prefetching		
Exam		
Break		
Caches + TLBs + Virtual memory + ChampSim framework	Talk 2 - Mainak Chaudhuri, IITK	
DRAM Controllers and DRAM address mapping		
Timing constraints, DRAM schedulling		
Timing constraints, DRAM schedulling, and DRAM cache		
Cache Coherence and memory consistency		

Processor [Source: Intel and Qualcomm]

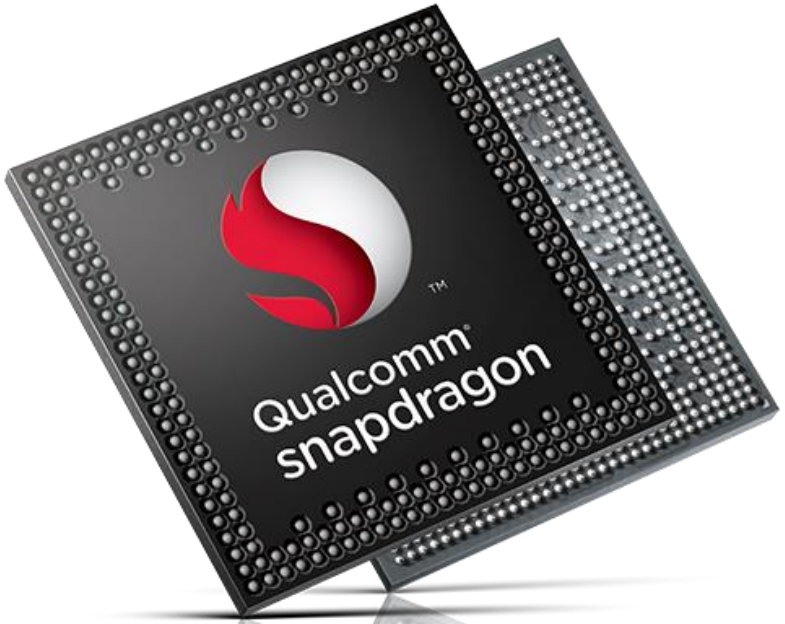

Caches [Source: Intel]

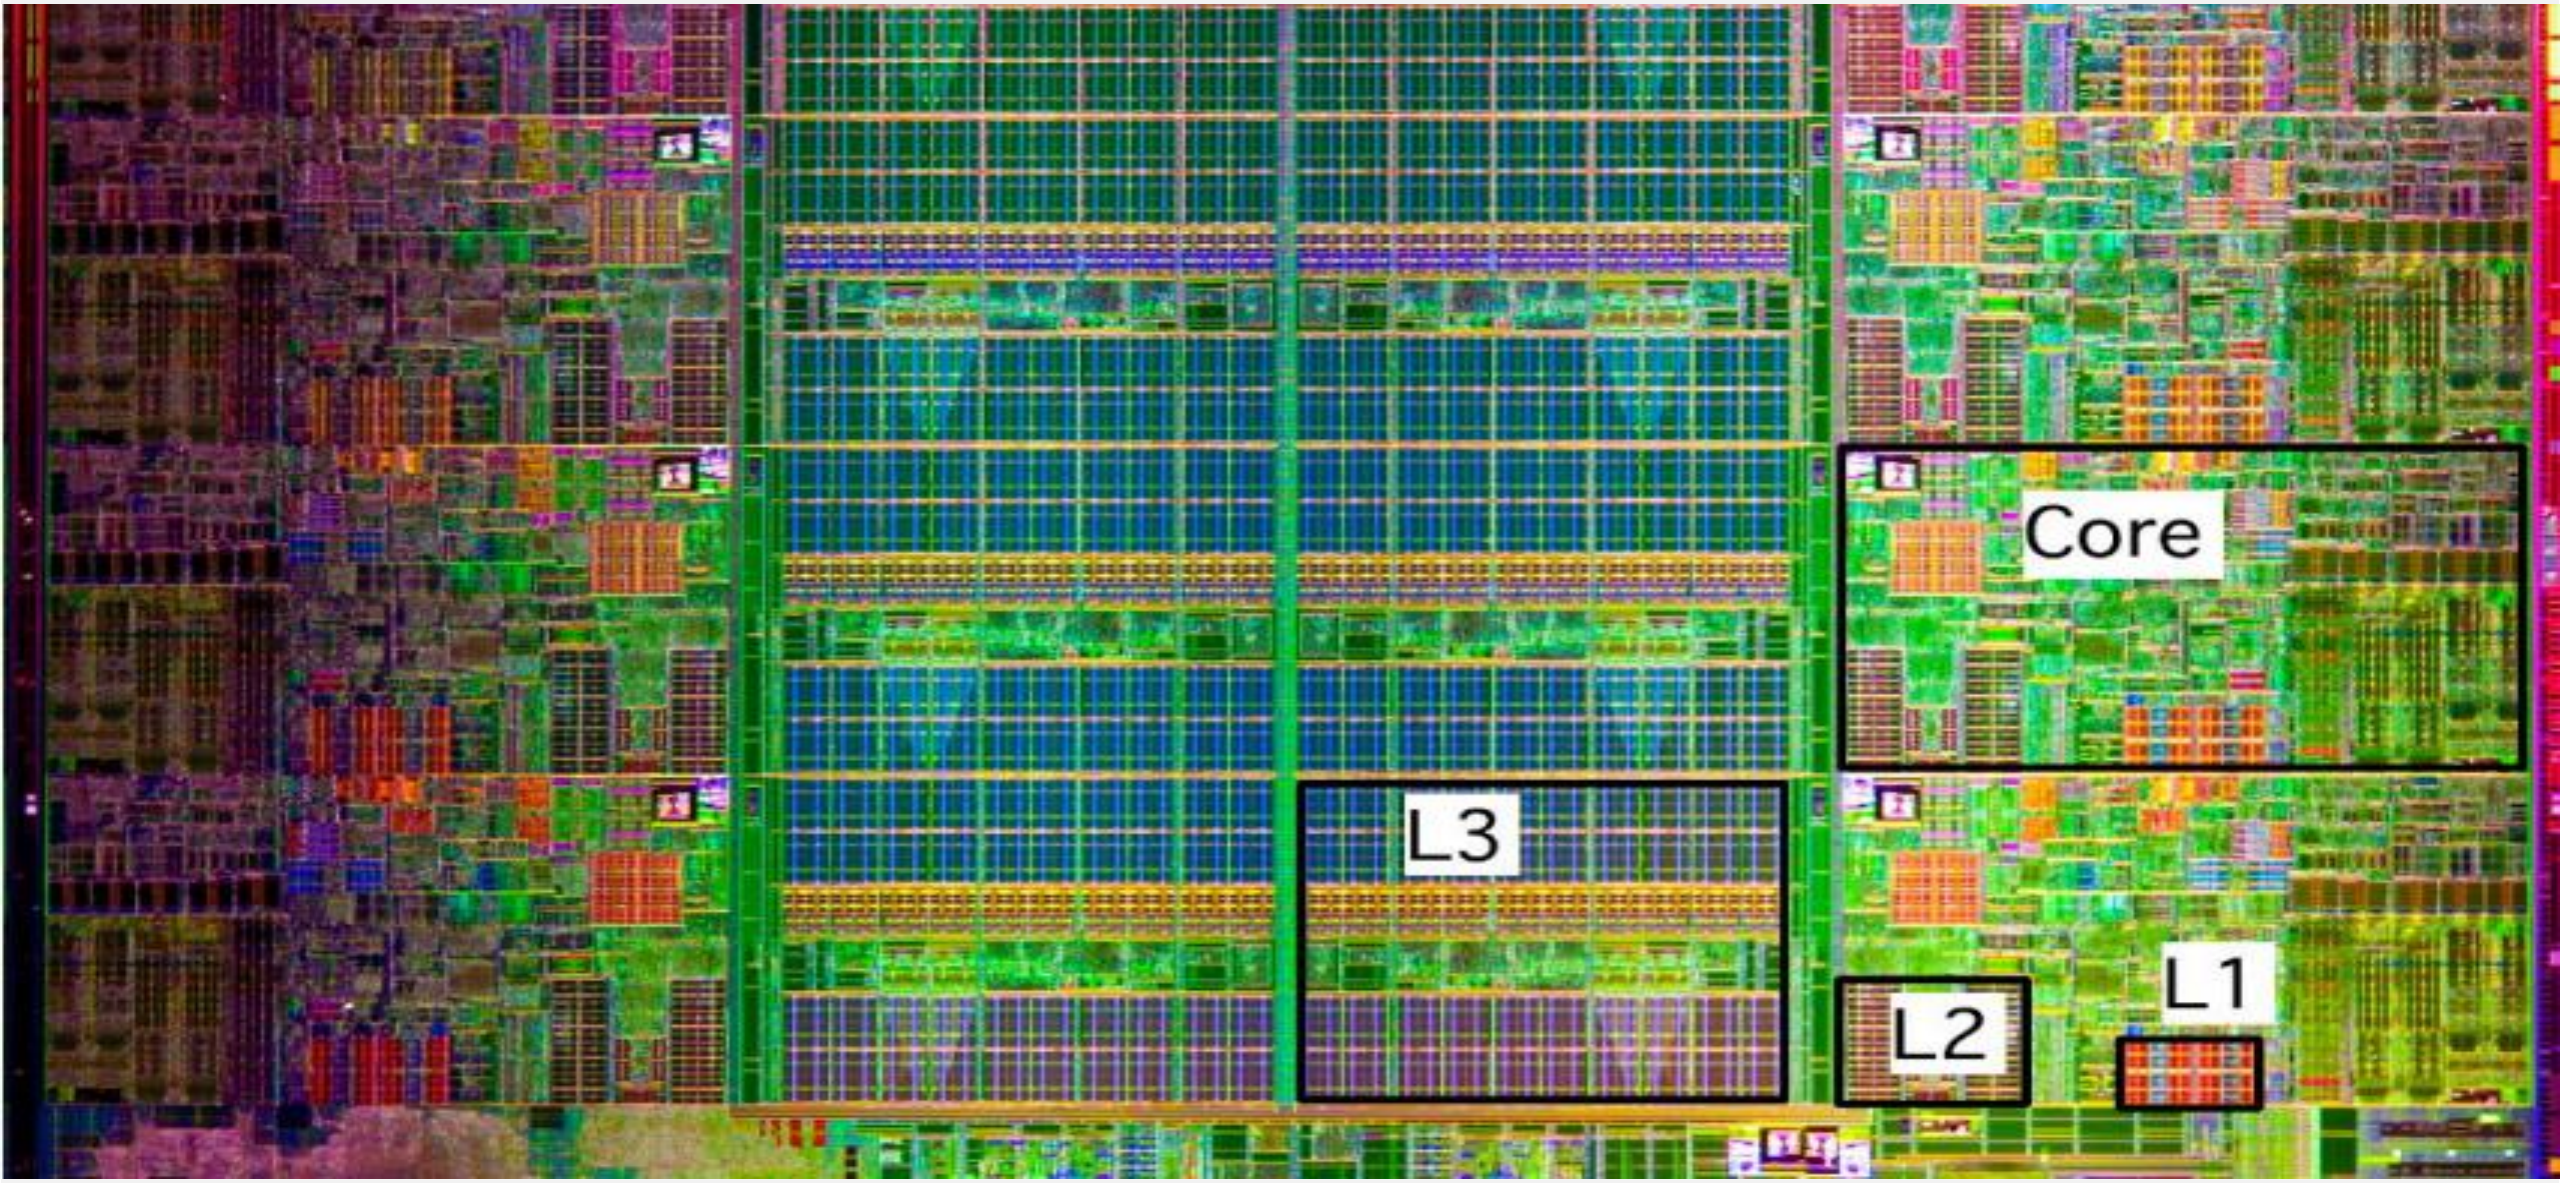

DRAM & DRAM Cache [Source: nist.gov]

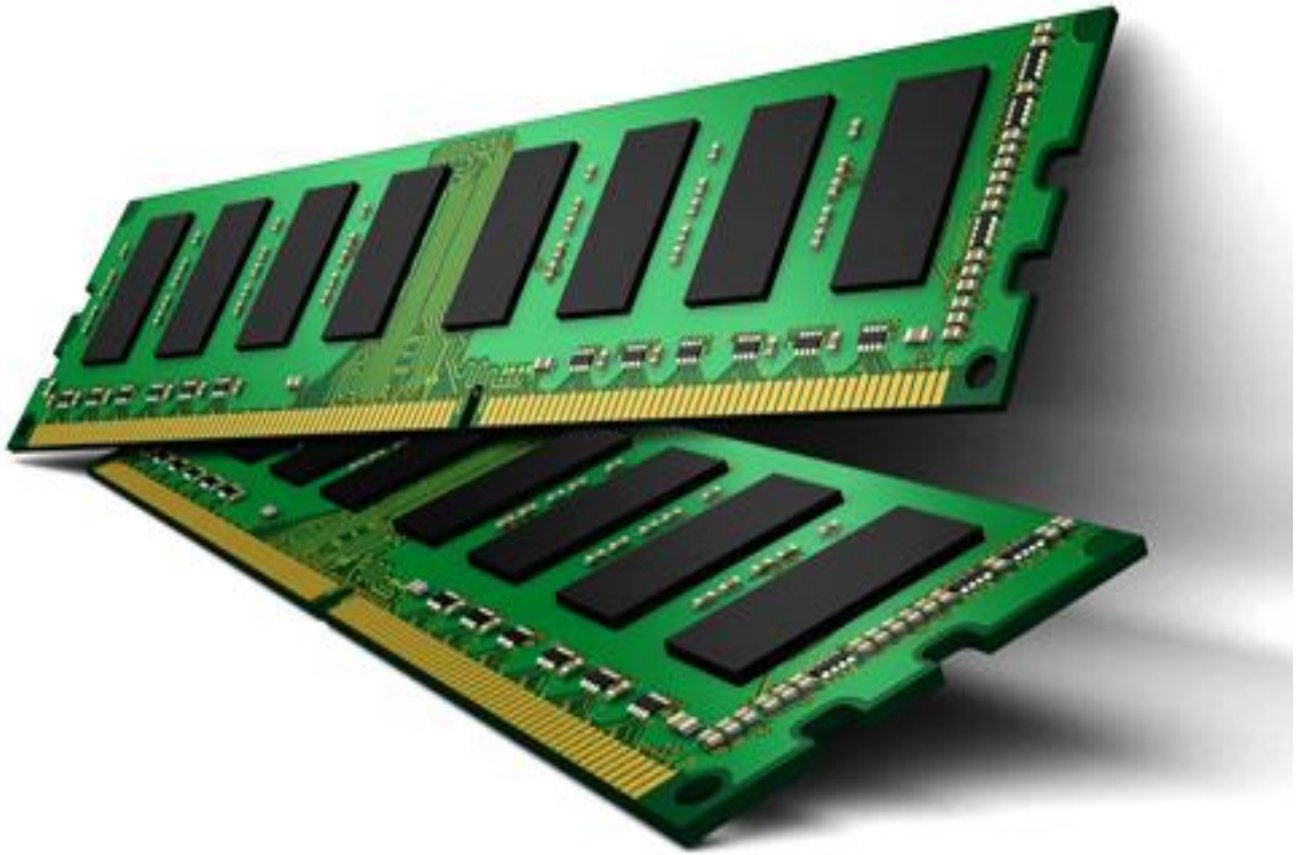

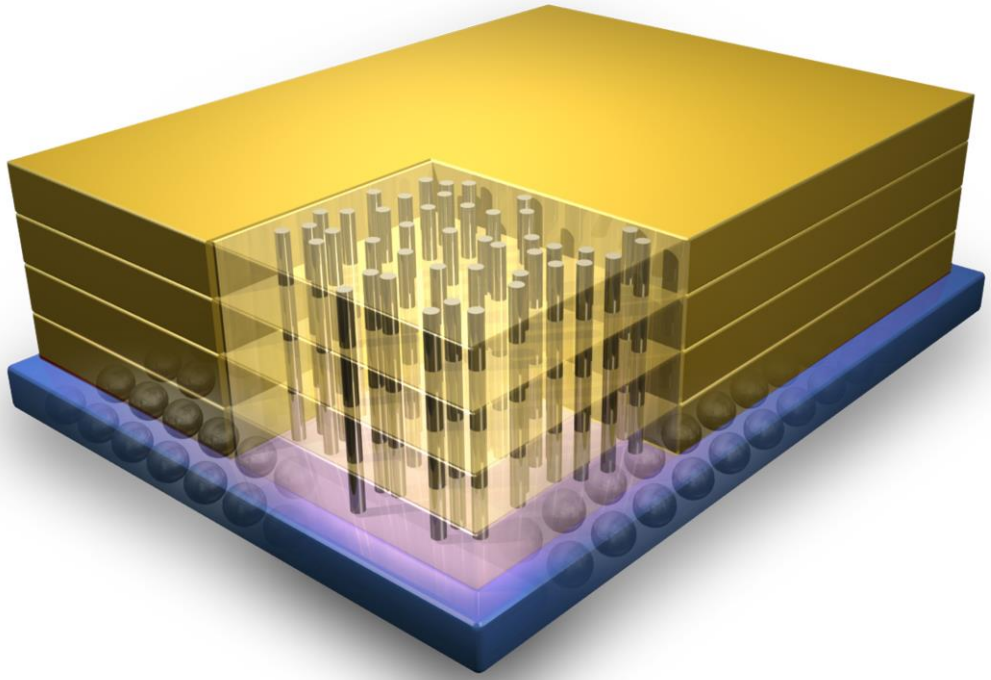

Others [Source: Youtube and NVIDIA]

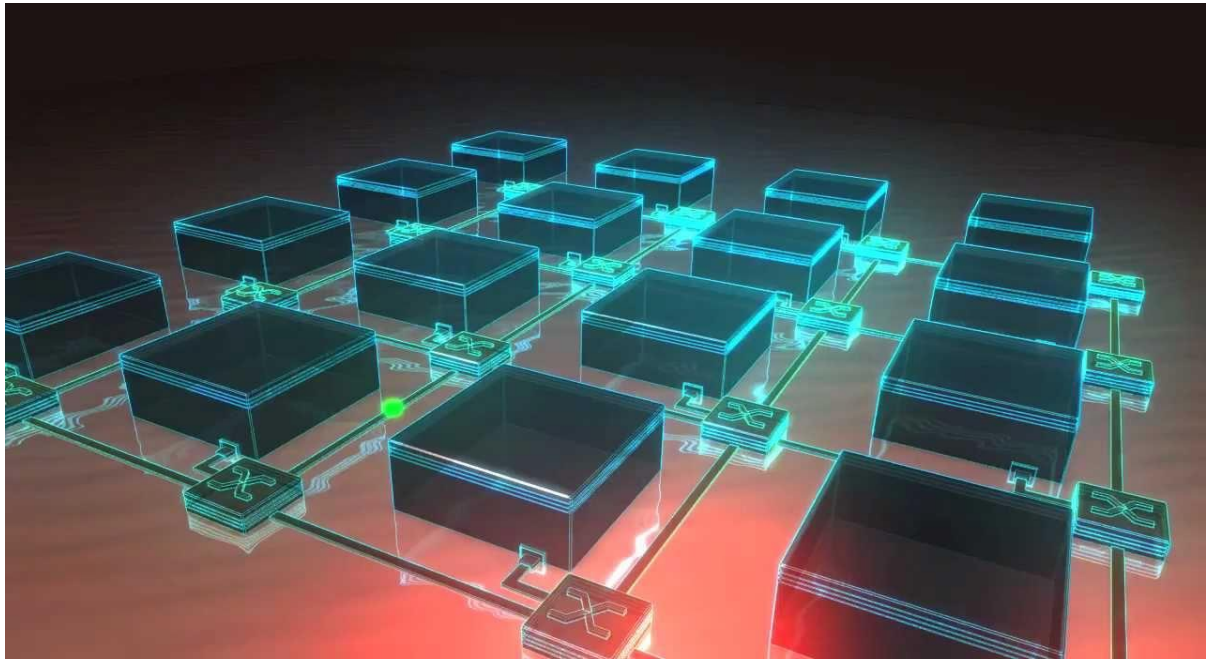

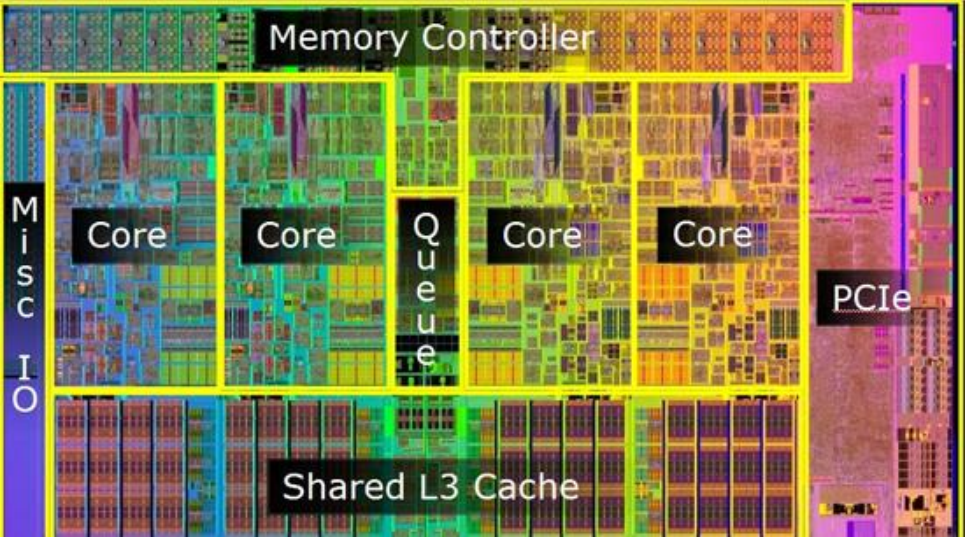

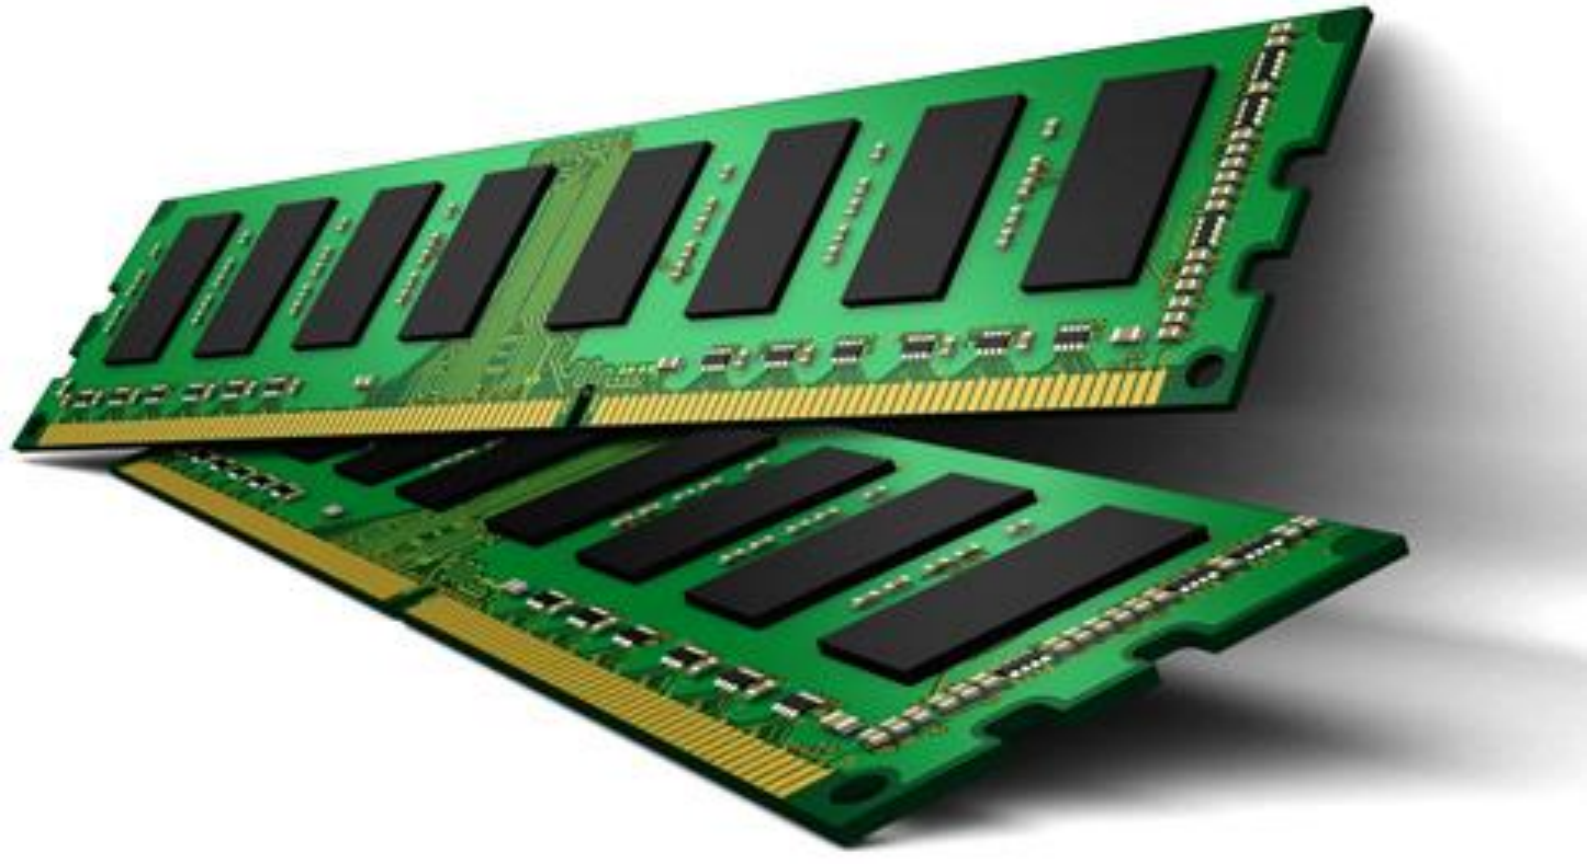

CS 422: Computer Architecture

Book

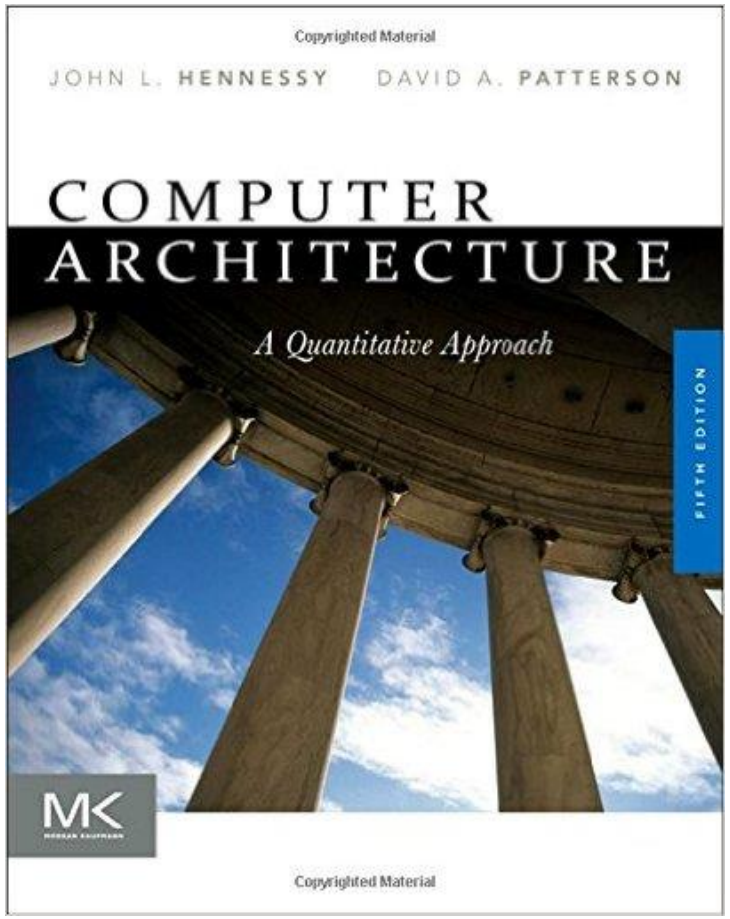
Copyrighted Material
JOHN L. HENNESSY DAVID A. PATTERSON
COMPUTER ARCHITECTURE
A Quantitative Approach
FIFTH EDITION
MK
Copyrighted Material

Let’s Stop for the SAFE

- Thanks to Bhaskar and Kameswari @CSE-IITB

Big News in Computer Architecture in Last Two Years

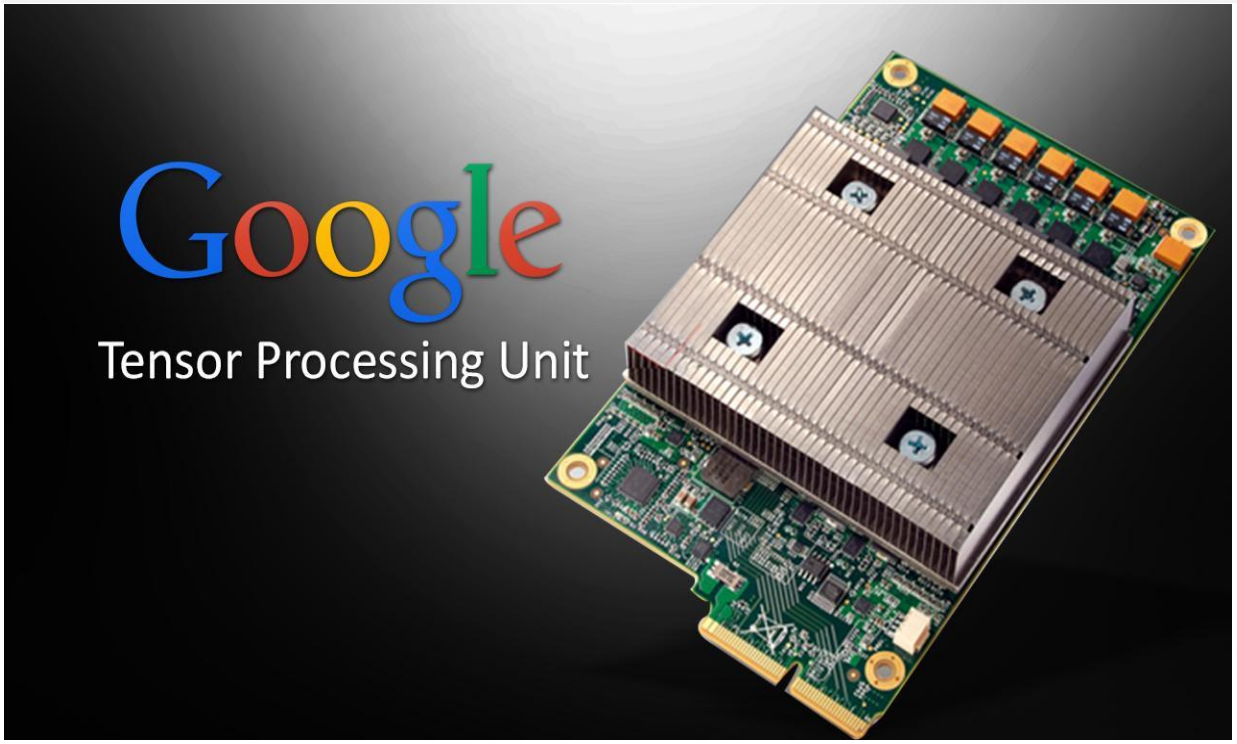

Source: Google

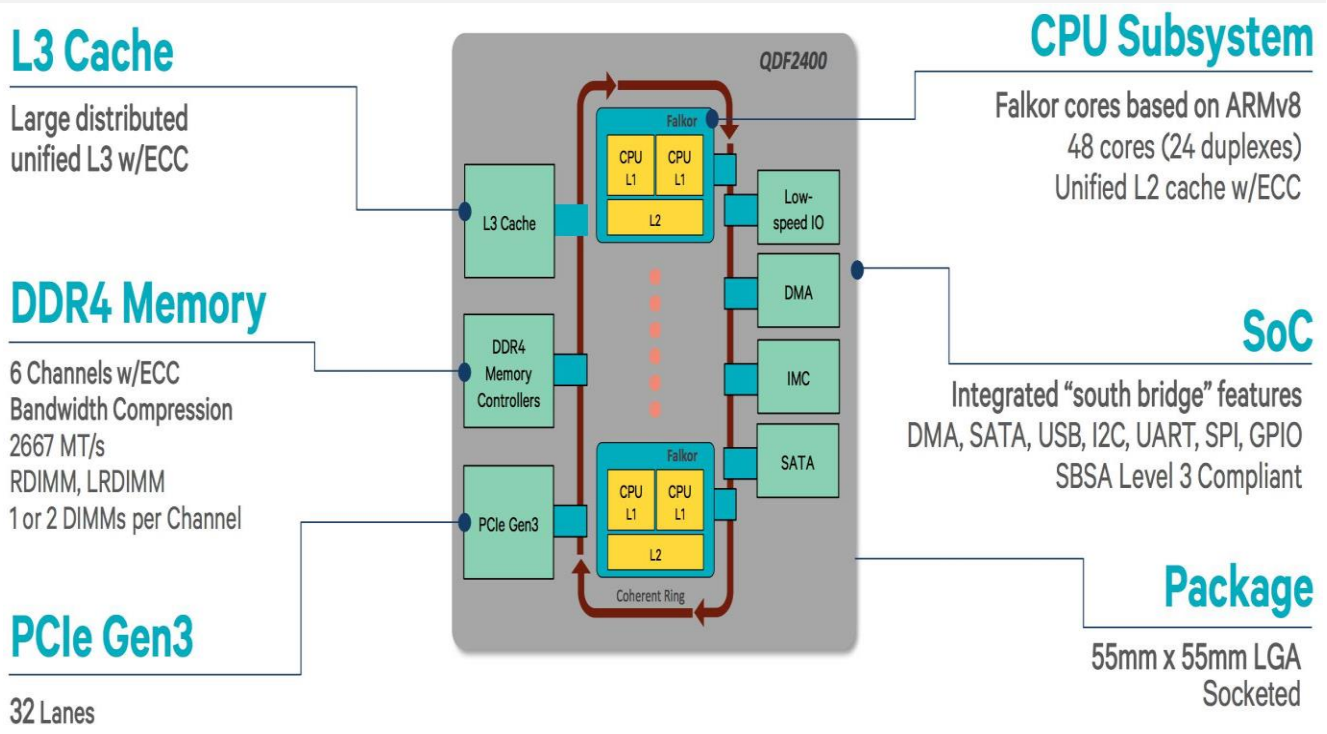

Source: Qualcomm

Brain Chip [Source: IBM]

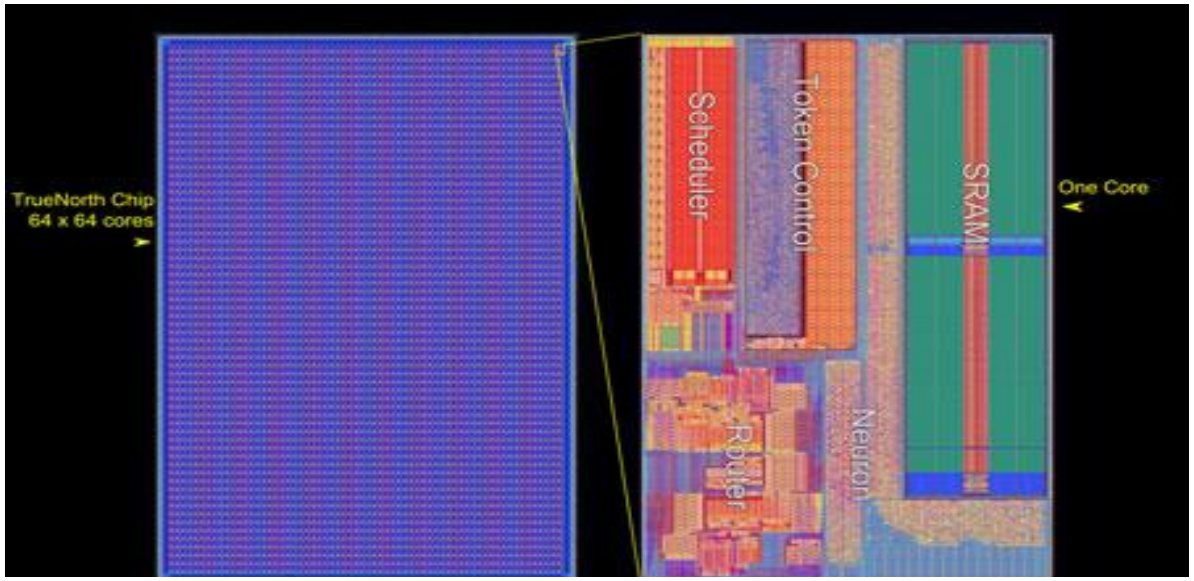

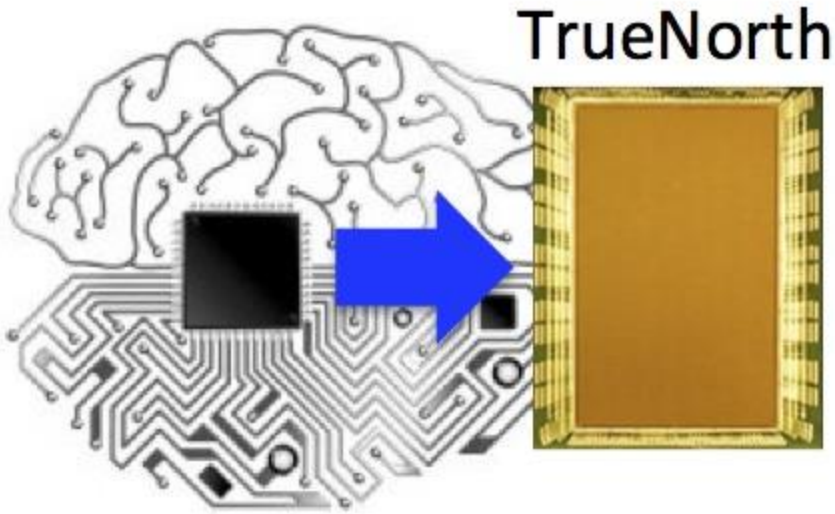

TrueNorth

1 M Neurons

256 M Synapses

Real time

73 mW

RISC – V [Source: UCB and Cadence]

Just Two Days Ago ☺

Processor flaw exposes 20 years of devices to new attack

Chipocalypse now

By Russell Brandom | @russellbrandom | Jan 3, 2018, 5:23pm EST

Meltdown

Meltdown breaks the most fundamental isolation between user applications and the operating system. This attack allows a program to access the memory, and thus also the secrets, of other programs and the operating system.

If your computer has a vulnerable processor and runs an unpatched operating system, it is not safe to work with sensitive information without the chance of leaking the information. This applies both to personal computers as well as cloud infrastructure. Luckily, there are software

Spectre

Spectre breaks the isolation between different applications. It allows an attacker to trick error-free programs, which follow best practices, into leaking their secrets. In fact, the safety checks of said best practices actually increase the attack surface and may make applications more susceptible to Spectre

Spectre is harder to exploit than Meltdown, but it is also harder to mitigate. However, it is possible to prevent specific known exploits based on Spectre through software patches.

Hang on! What is Computer Architecture?

Computer Architecture ??

Since 1946 all computers have had 5 components

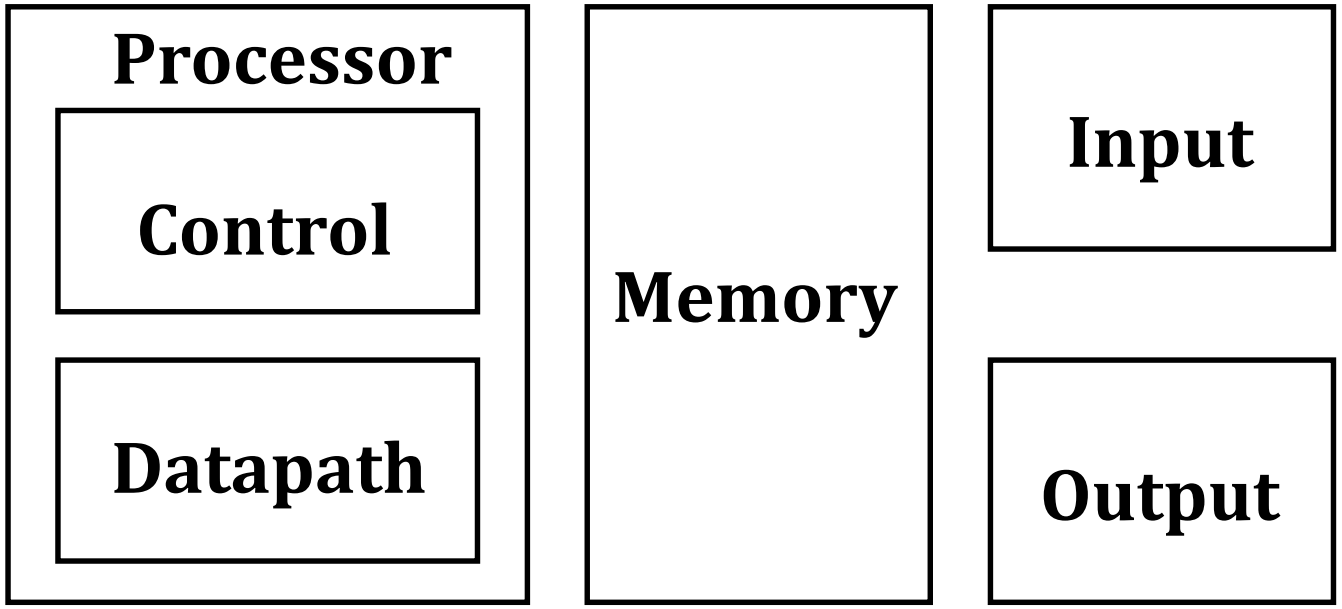

So What ?

Remember CS220

5-stage instruction pipeline

RISC/CISC

Caches

ISA

Addressing modes

CPI, IPC, MIPS Rating

Again, What is Computer Architecture?

- 1950s to 1960s: Computer Arithmetic
- 1970s to mid 1980s: Instruction Set Design, especially ISA appropriate for compilers
- 1990s: Design of CPU, memory system, I/O system, Multiprocessors, Networks
- 2010s: Self adapting systems? Self organizing structures? DNA Systems/Quantum Computing?

What is Computer Architecture?

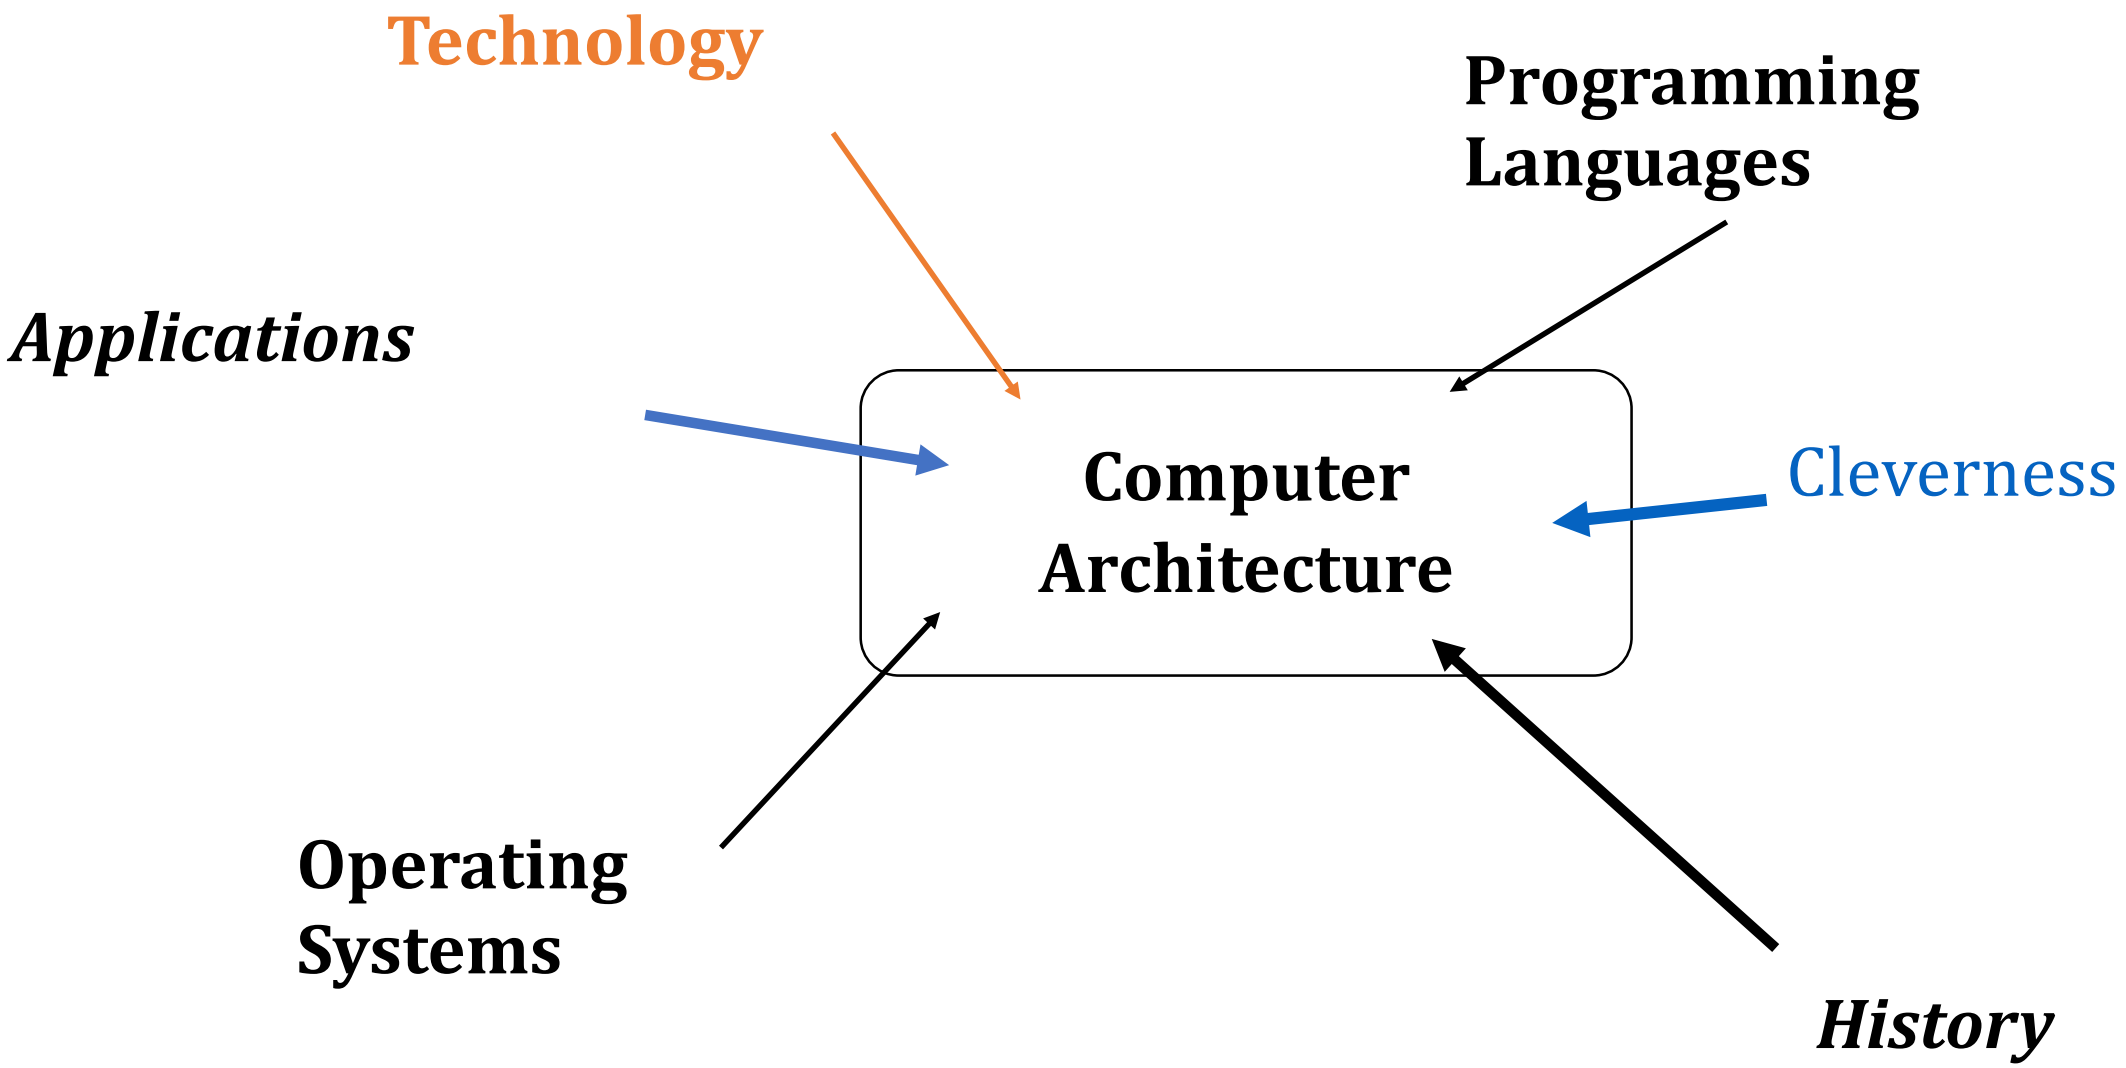

Computer Architecture ??

VLSI++ or ++VLSI
Writing Verilog/VHDL code for designing a processor
Understanding how transistors work

Computer theorists propose algorithms that solve important problems and analyze their asymptotic behavior (e.g., O(NlogN), O(N)). Computer architects (applicable to computer systems) set the constant factors of these algorithms – Christos Kozyrakis, Stanford

Computer Architecture: My View

For non-CS/EE minds: Abstraction layer that enables computation in (running a C program and getting an output) hardware. The layer decides how/when/why of the enabler.

For CS/EE minds: Study of design trade-offs of different components (five) that are part of the abstraction layer. Trade-offs can be in terms of performance, power, energy, area, security, ……..

Computing Stack

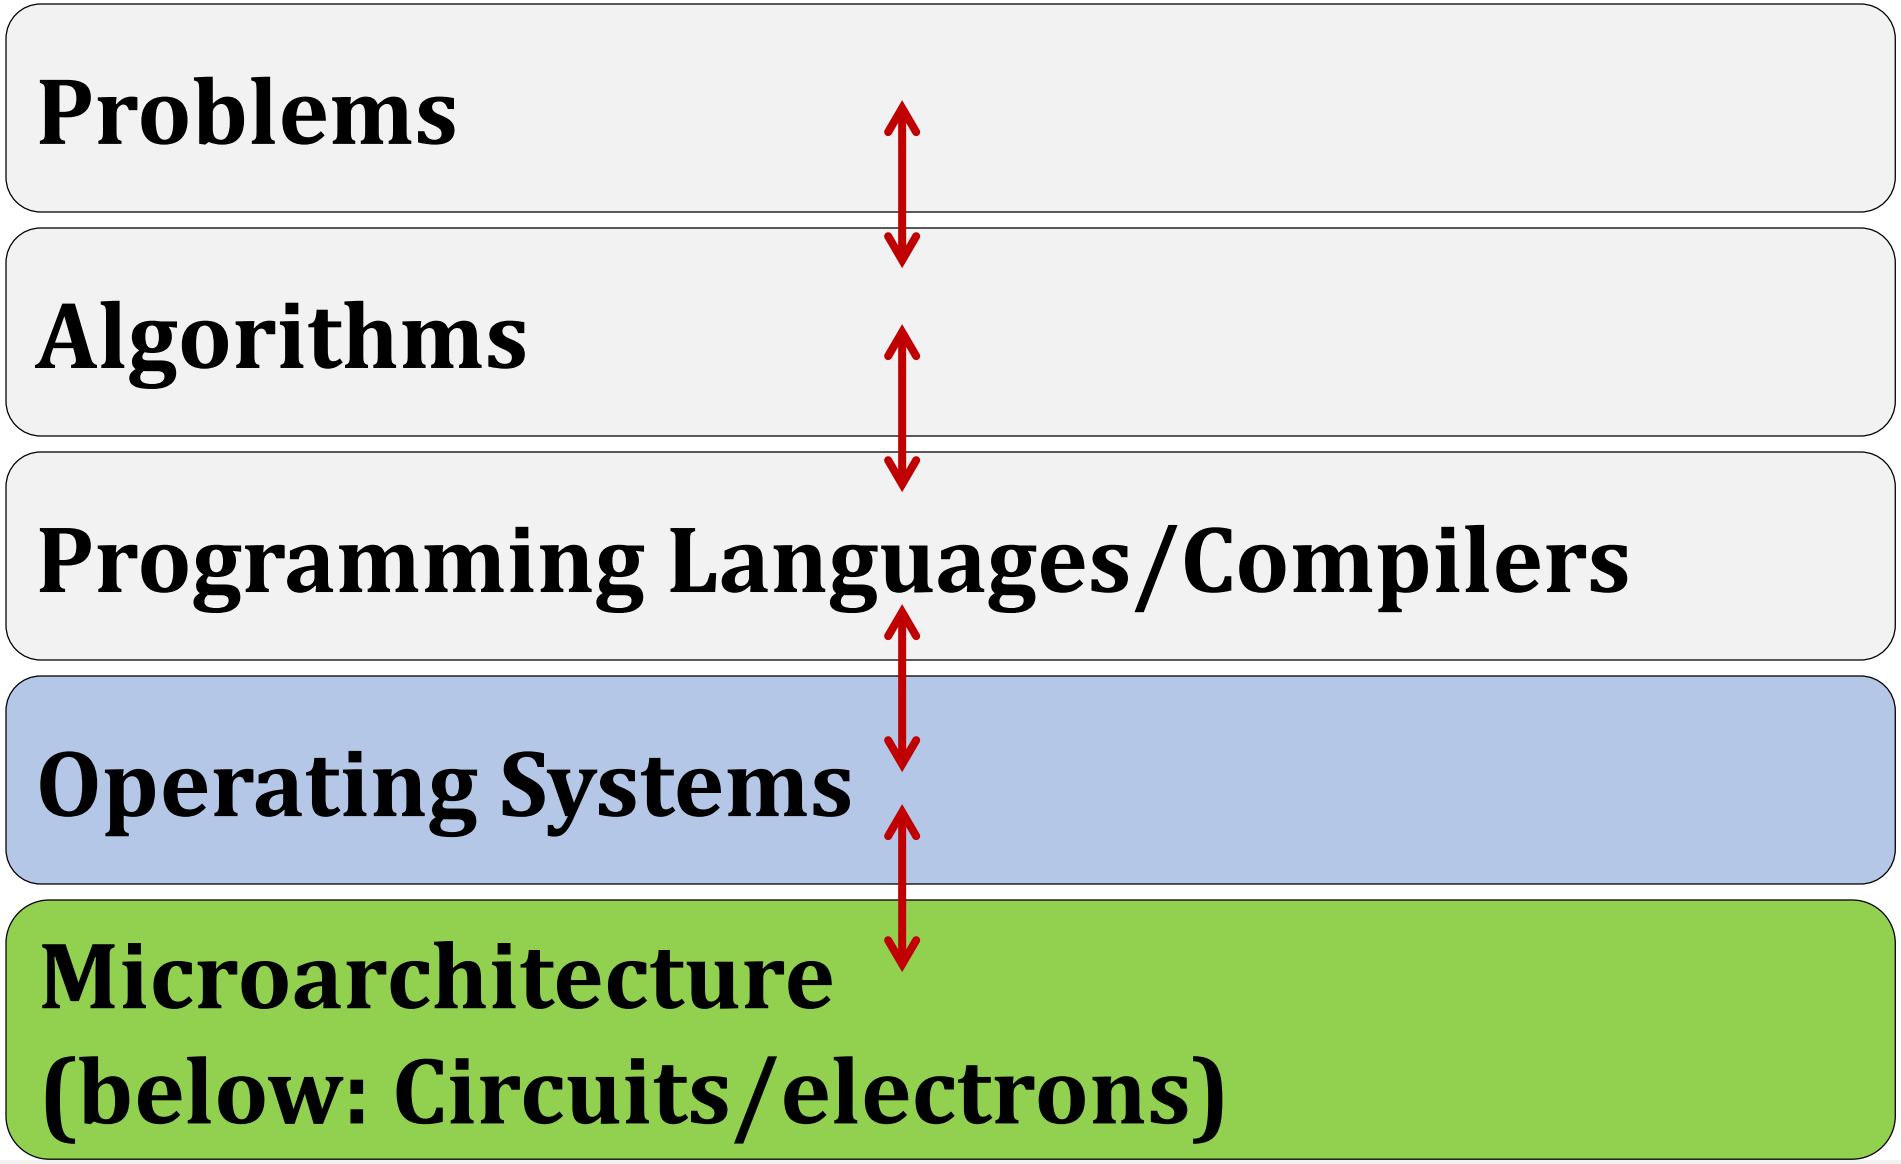

Why Study Computer Architecture ?

- CHANGE
- It's exciting!
- It has never been more exciting!
- It impacts every other aspect of electrical engineering and computer science

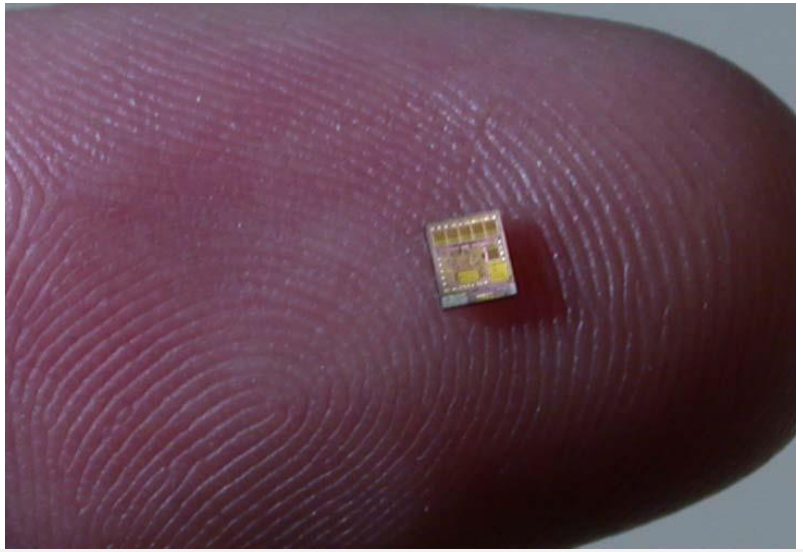

Architecture: Design Process

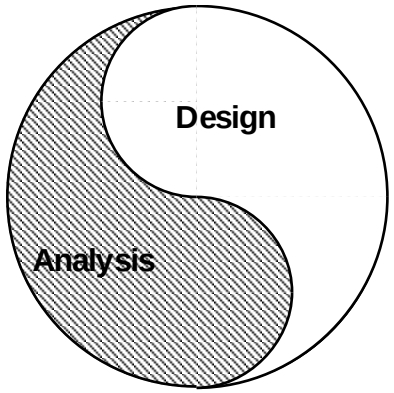

Architecture is an iterative process:

- Searching the space of possible designs
- At all levels of computer systems

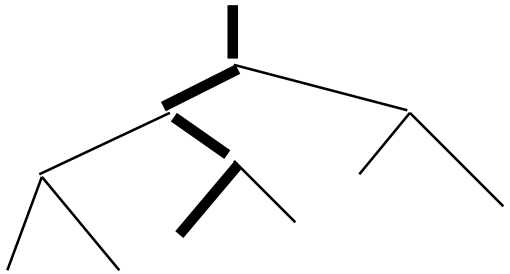

Creativity

Cost /Performance Analysis

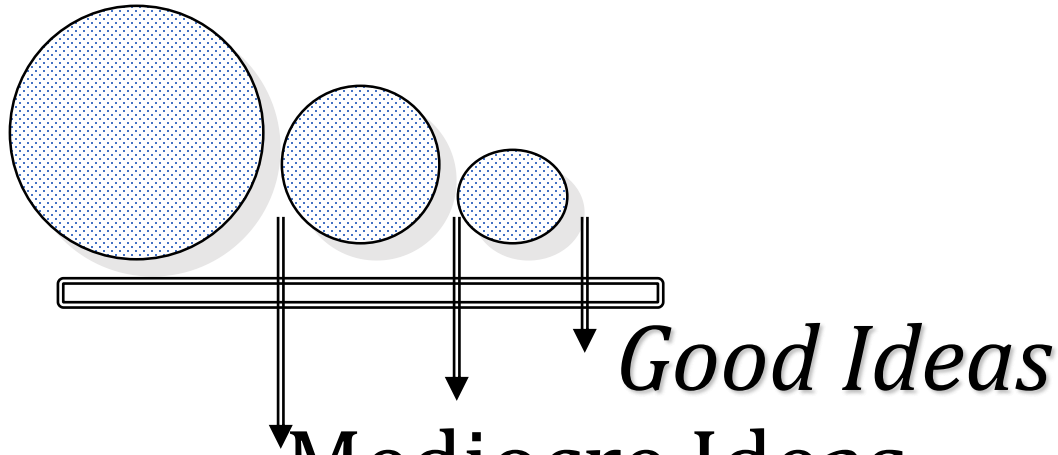

First Microprocessor: Intel 4004, 1971

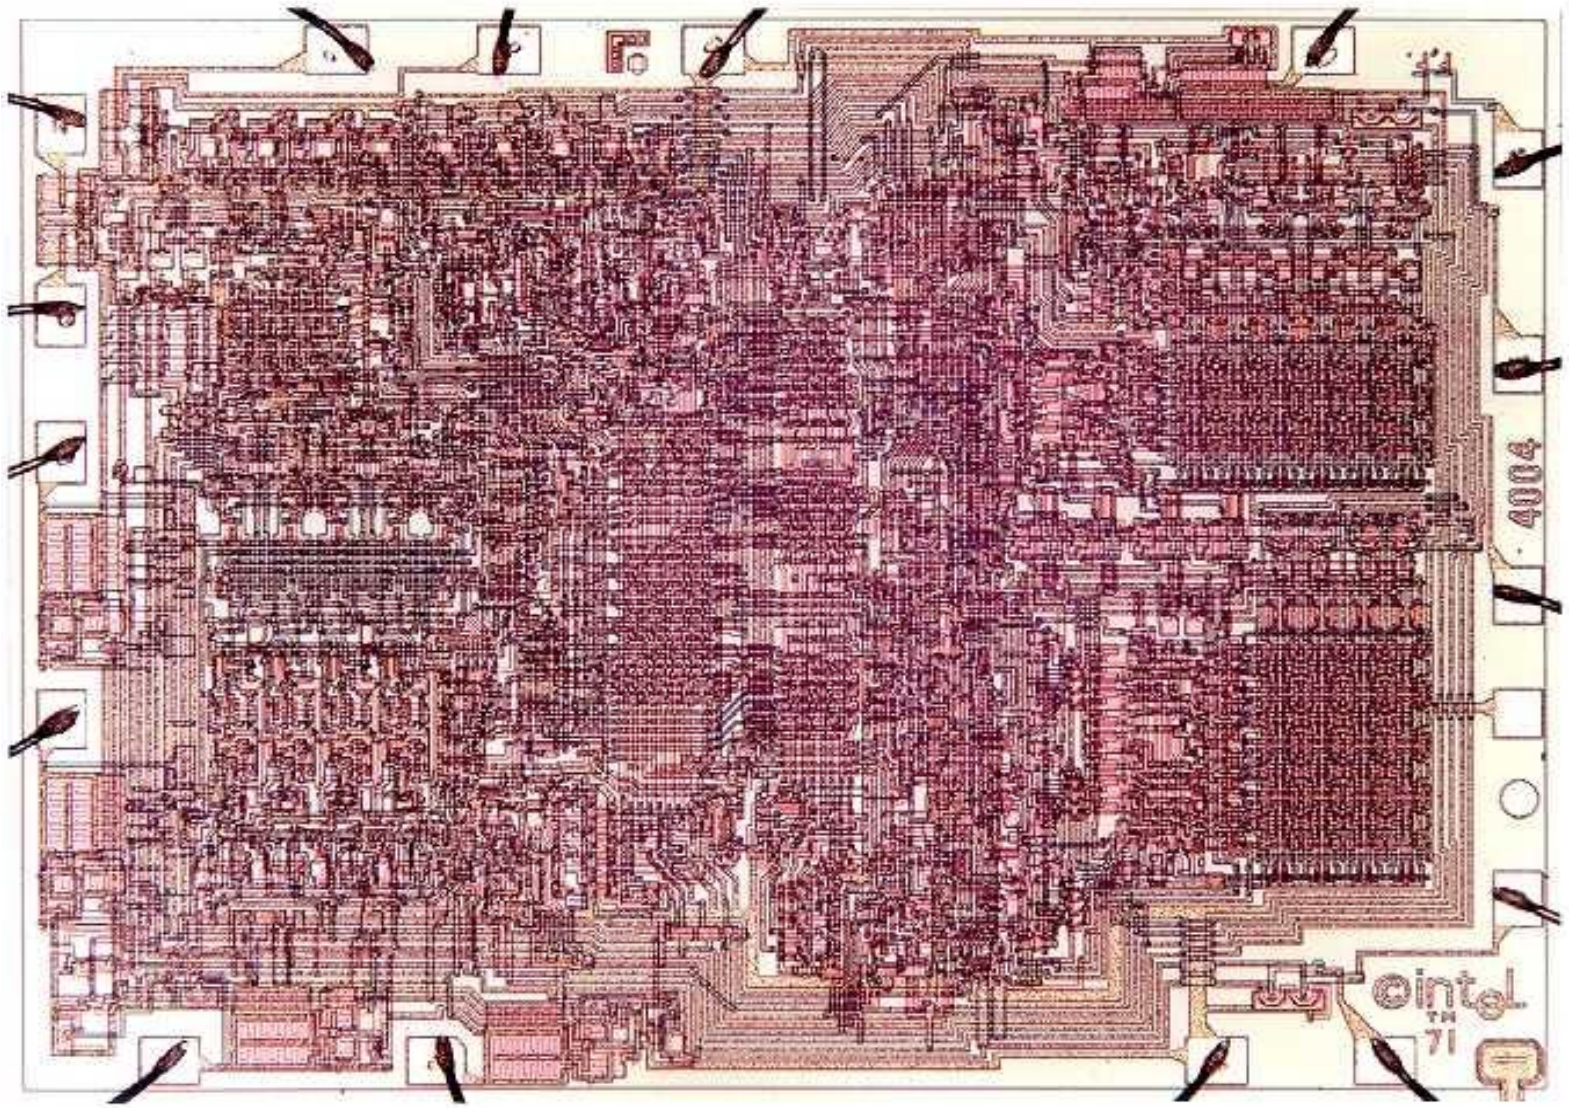

- 4-bit accumulator architecture
- 8µm pMOS
- 2,300 transistors
- 3 x 4 mm2
- 750kHz clock
- 8-16 cycles/inst.

Before That: Difference Engine

1855. Can compute any 6th degree polynomial by calculating the difference between 2D matrix elements

Speed: 33 to 44 32-digit numbers per minute!

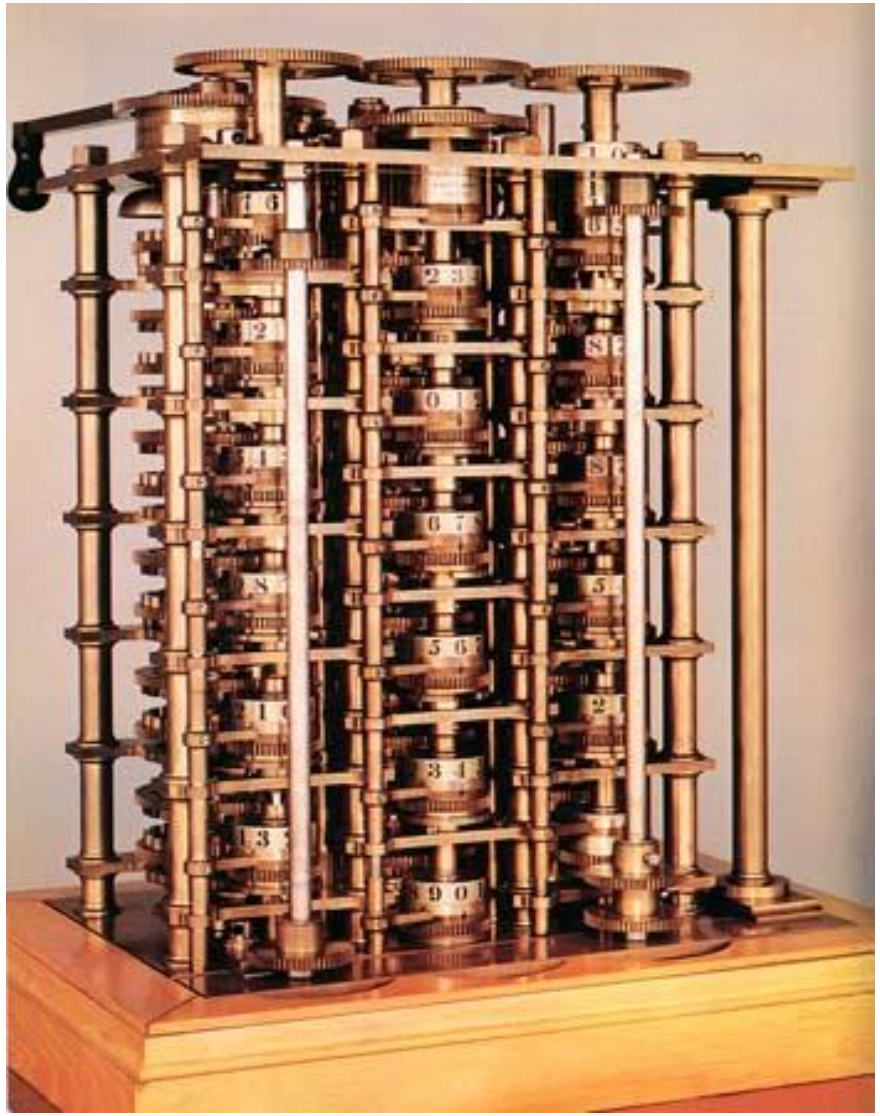

Harvard Mark -I

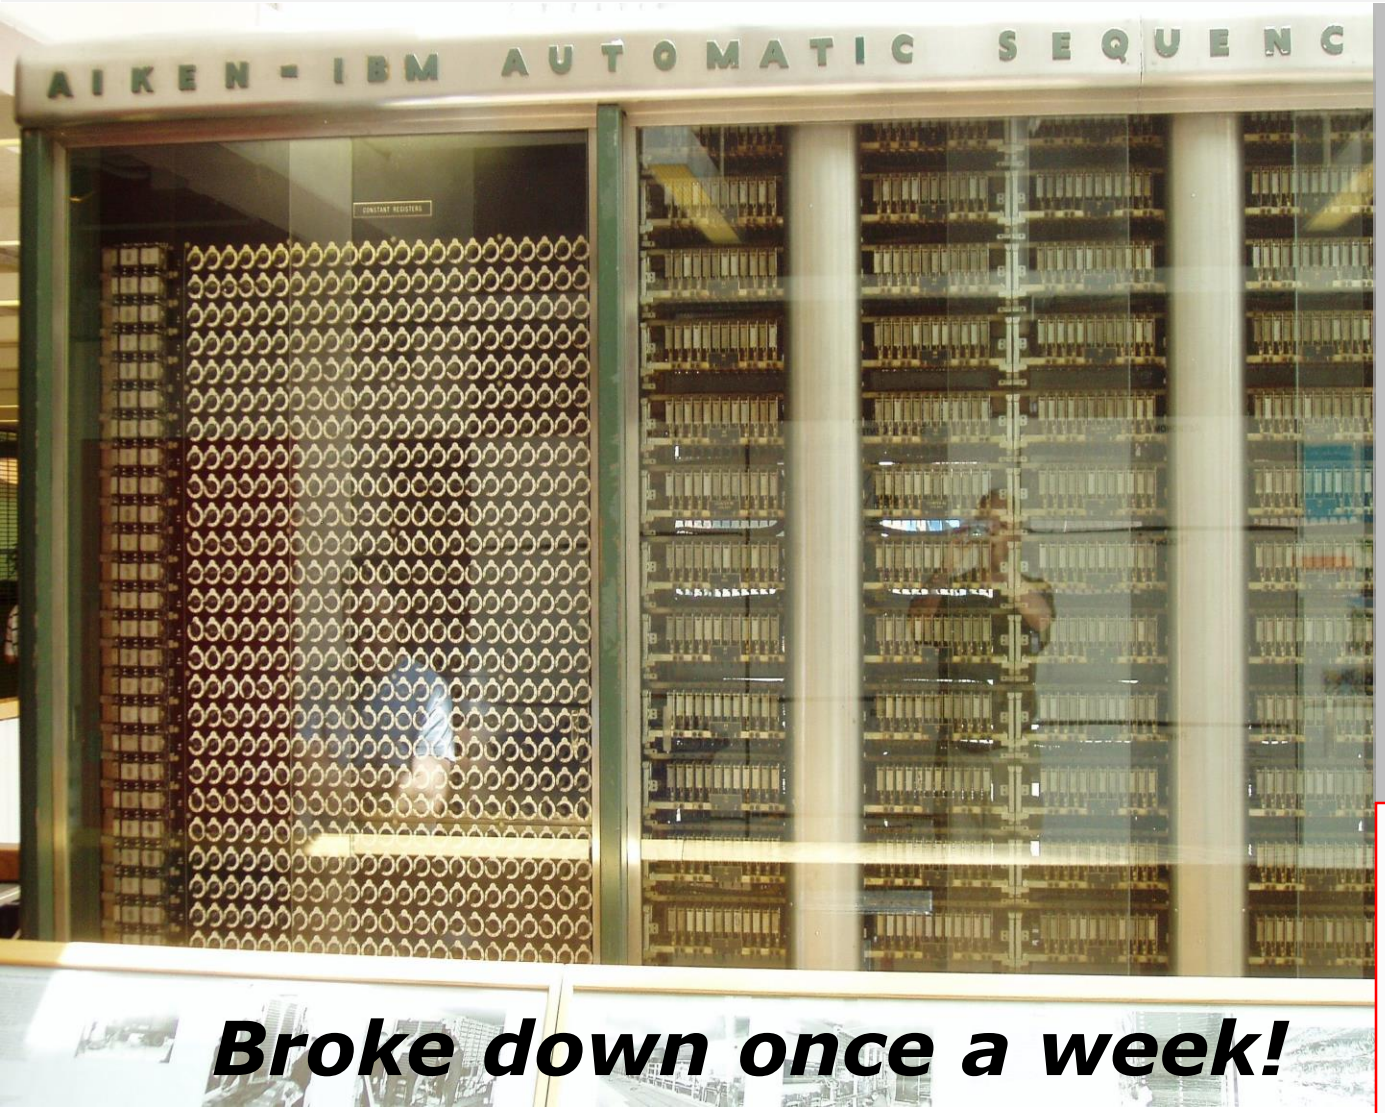

Broke down once a week!

- **Built in 1944 in IBM Endicott laboratories**
  - **Howard Aiken – Professor of Physics at Harvard**
  - **Essentially mechanical**
  - **Weighed *5 tons* and had *750,000* components**
  - **A synchronizing clock that beat every *0.015* seconds (66Hz)**
  - **Inspired by Charles Babbage's analytic engine**

Performance:

0.3 seconds for addition
6 seconds for multiplication
1 minute for a sine calculation

EDSAC in 1949 (EDVAC in 1944 By V. Newmann)

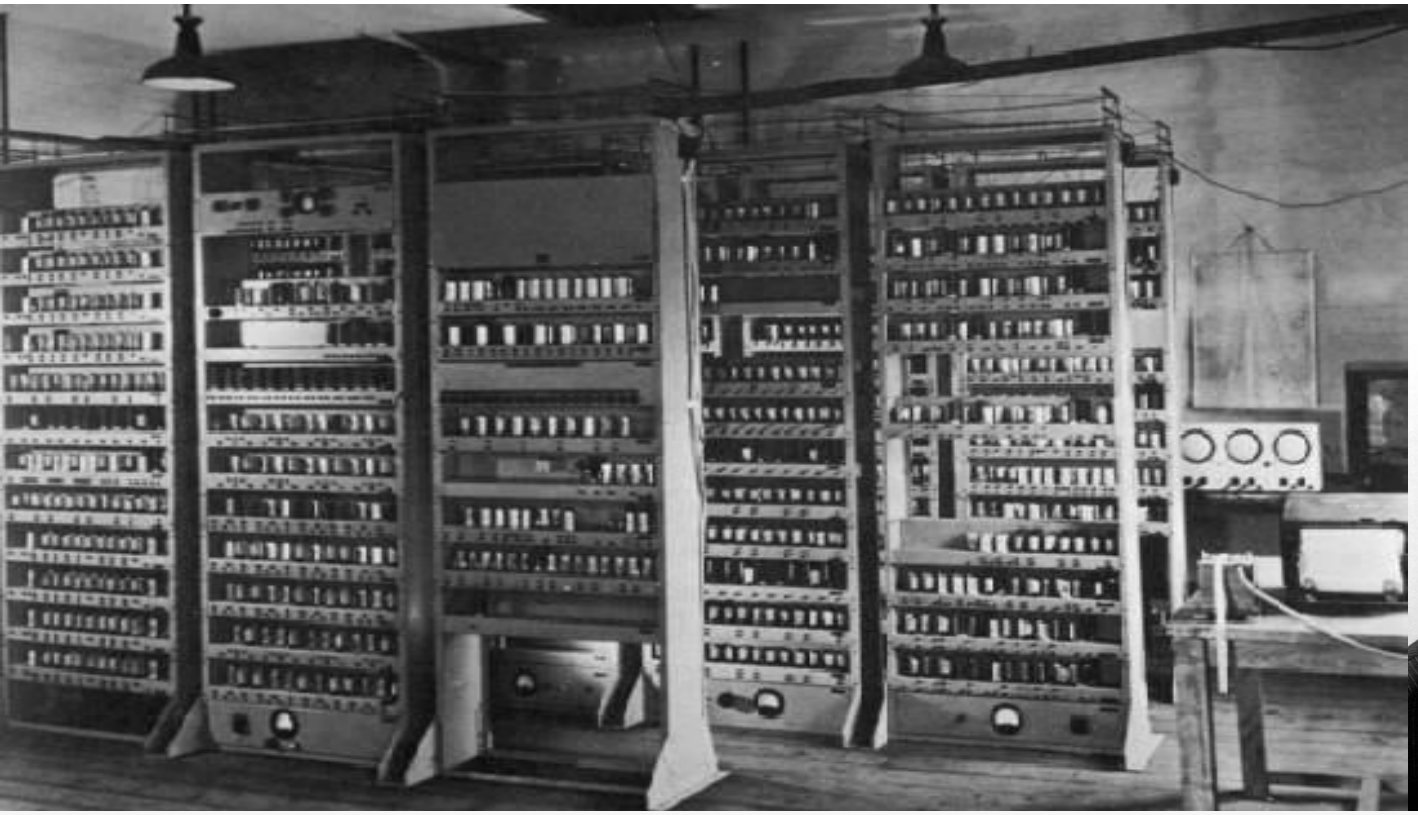

Electronic Delay Storage Automatic Calculator by Maurice Wilkes

Source: U. Cambridge

ACM SIGARCH Maurice Wilkes Award

The award of $2,500 is given annually for an outstanding contribution to computer architecture made by an individual whose computer-related professional career (graduate school or full-time employment, whichever began first) started no earlier than January 1st of the year that is 20 years prior to the year of the award.*

IBM 701

IBM 701 (1952)

30 machines sold by IBM in 1953-54

Why IBM entered so late into ..?

IBM revenues were doubling every 4 to 5 years in 1940/50s. Remember, this is without computers.

Intel's 8086

1978: Around 50 instructions

2017: Around 650 instructions

Personal Computing

Presenting the IBM® of Personal Computers.

IBM is proud to announce a product *you* may have a personal interest in. It's a tool that could soon be on your desk, in your home or in your child's schoolroom. It can make a surprising difference in the way you work, learn or otherwise approach the complexities (and some of the simple pleasures) of living.

It's the computer we're making for you.

In the past 30 years, the computer has become faster, smaller, less complicated and less expensive. And IBM has contributed heavily to that evolution.

Today, we've applied what we know to a new product we believe in: the IBM Personal Computer.

IBM PERSONAL COMPUTER SPECIFICATIONS
*ADVANCED FEATURES FOR PERSONAL COMPUTERS

User Memory
16K - 256K bytes*
Permanent Memory
(ROM) 40K bytes*
Microprocessor
High speed, 8088*
Auxiliary Memory
2 optional internal diskette drives, 5¼", 160K bytes per diskette
Keyboard
83 keys, 6 ft. cord attaches to system unit*
10 funcuon keys*
10-key numeric pad
Tactile feedback*

Display Screen
High-resolution (720h x 350v)*
80 characters x 25 lines
Upper and lower case
Green phosphor screen*
Diagnostics
Power-on self testing*
Parity checking
Languages
BASIC, Pascal
Printer
Bidirectional*
80 characters/second
12 character styles, up to 132 characters/line*
9 x 9 character matrix*

Color/Graphics
Text mode:
16 colors*
256 characters and symbols in ROM*
Graphics mode:
4-color resolution: 320h x 200v*
Black & white resolution: 640h x 200v*
Simultaneous graphics & text capability*
Communications
RS-232-C interface
Asynchronous (start/stop) protocol
Up to 9600 bits per second

It's a computer that has reached a truly personal scale in size and in price: starting at less than $1,600† for a system that, with the addition of one simple device, hooks up to your home TV and uses your audio cassette recorder.

For flexibility, performance and ease of use, no other personal computer offers as many advanced features to please novice and expert alike (see the box).

Features like high resolution color graphics. Ten, user-defined function keys. The kind of expandability that lets you add a printer for word processing, or user memory up to 256KB. Or BASIC and Pascal languages that let you write your own programs. And a growing list of superior programs like VisiCalc,™ selected by IBM to match the quality and thoughtfulness of the system's total design.

This new system will be sold through channels which meet our professional criteria: the nationwide chain of 150 ComputerLand® stores, and Sears Business Systems Centers. Of course, our own IBM Product Centers will sell and service the system. And the IBM Data Processing Division will serve those customers who want to purchase in quantity.

Experience the IBM Personal Computer. You'll be surprised how quickly you feel comfortable with it. And impressed with what it can do for you.

IBM®

For the IBM Personal Computer dealer nearest you, call (800) 447-4700. In Illinois, (800) 322-4400.

†This price applies to IBM Product Centers. Prices may vary at other stores.
VisiCalc is a trademark of Personal Software, Inc.

CIRCLE 3

[Personal Computing Ad, 11/81]

Do You Know What It Is ?

Today: Datacenter @Google

Today: China's Sunway Taihulight

(more than 100K cores)

And We are Smart now ☺

Source:
Truthseeker, UK

Hang on !! Why All This ?

You have to understand the past to understand the present: Carl Sagan

Let’s Look at the Other (Consumer) Side

World of Programming Languages

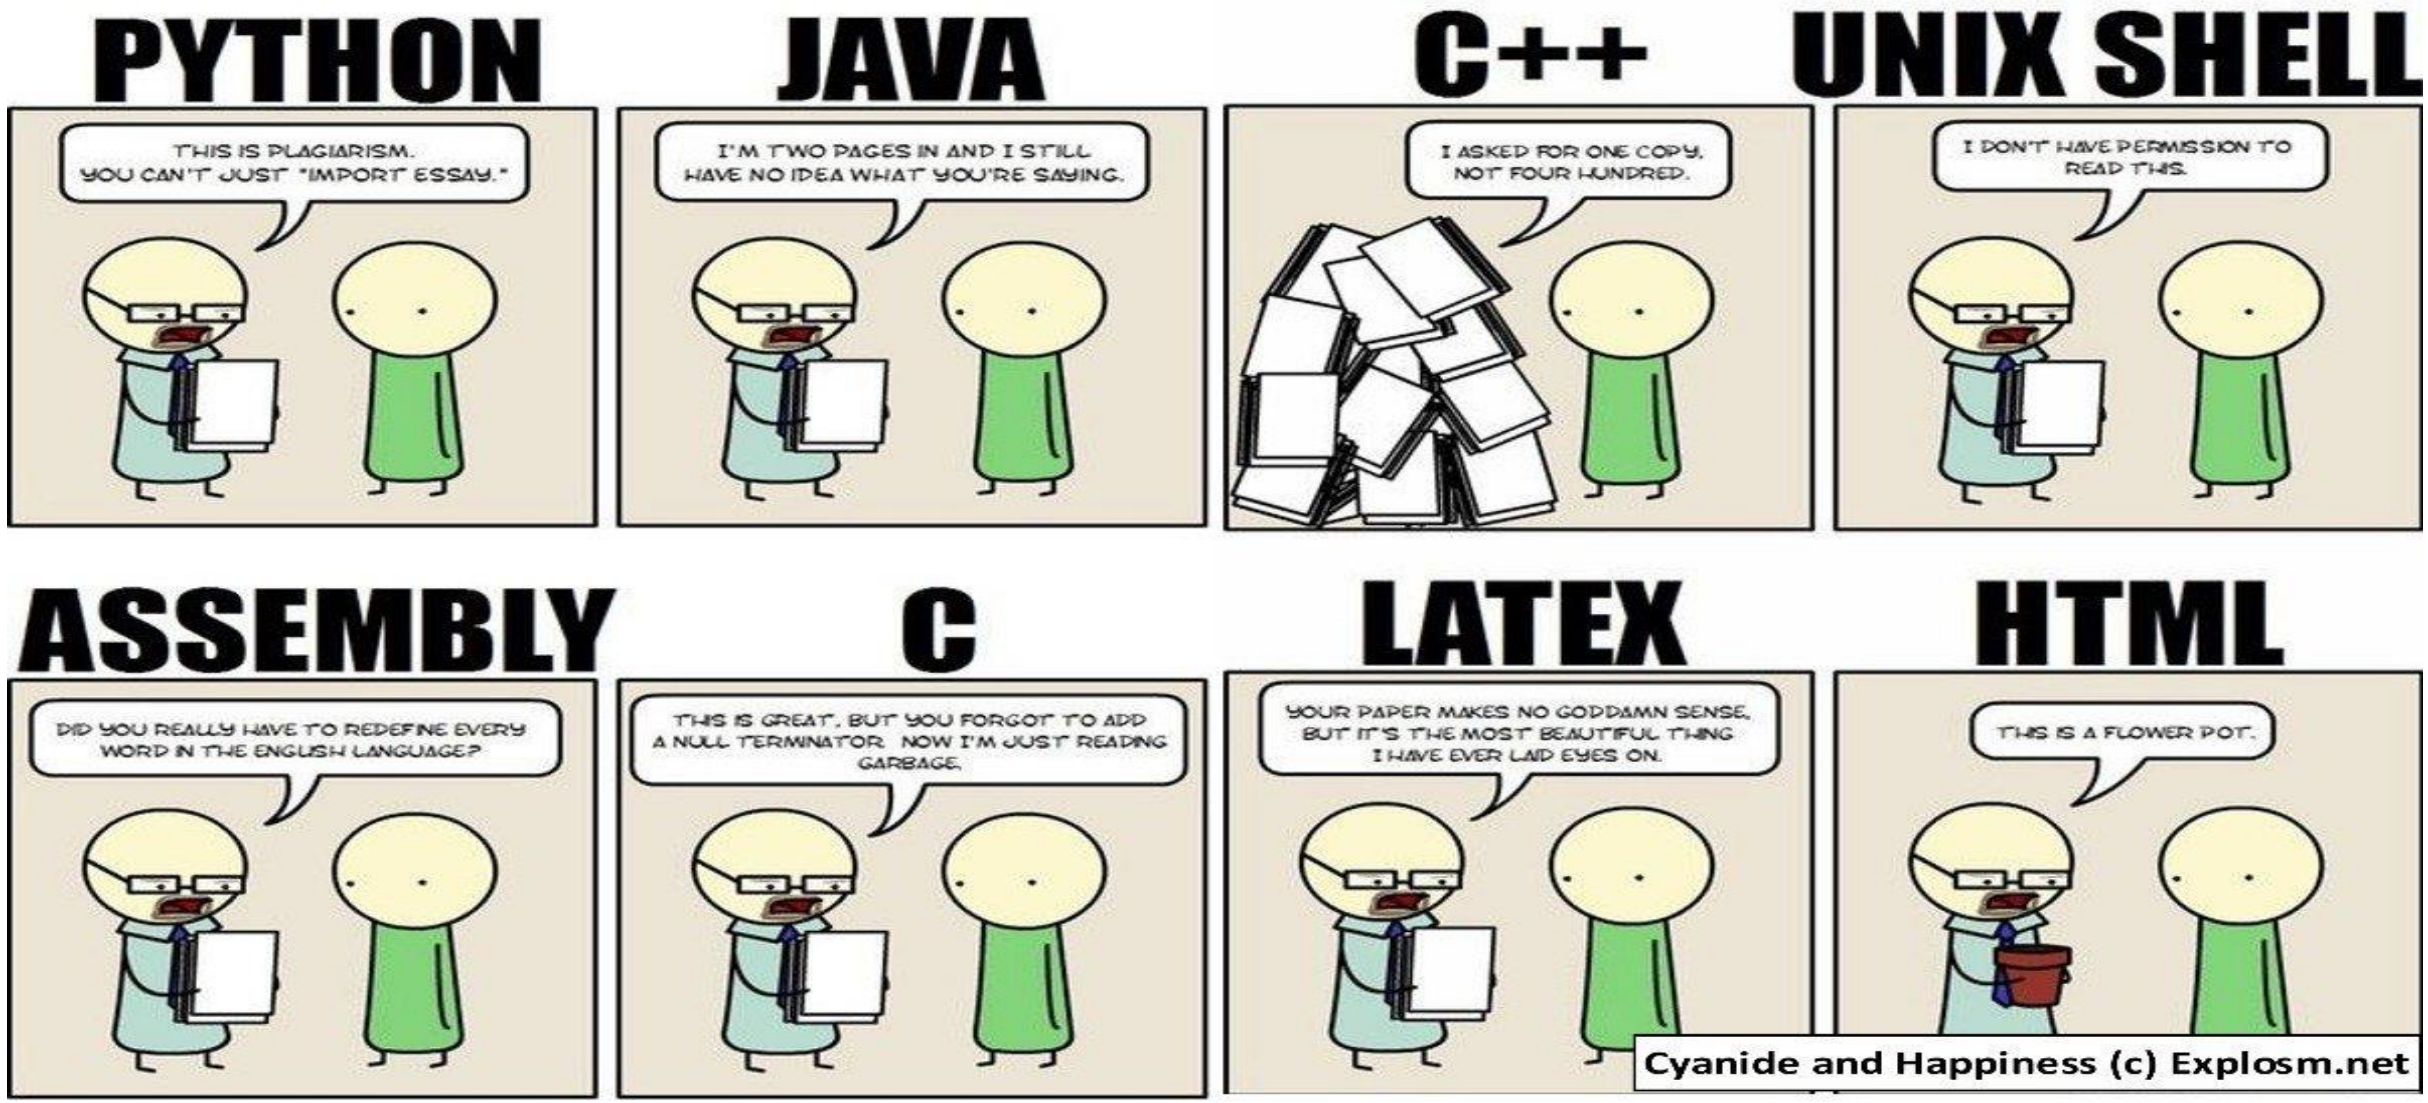

What About Application domains ?

Look around you and think of it

Let's Revisit This

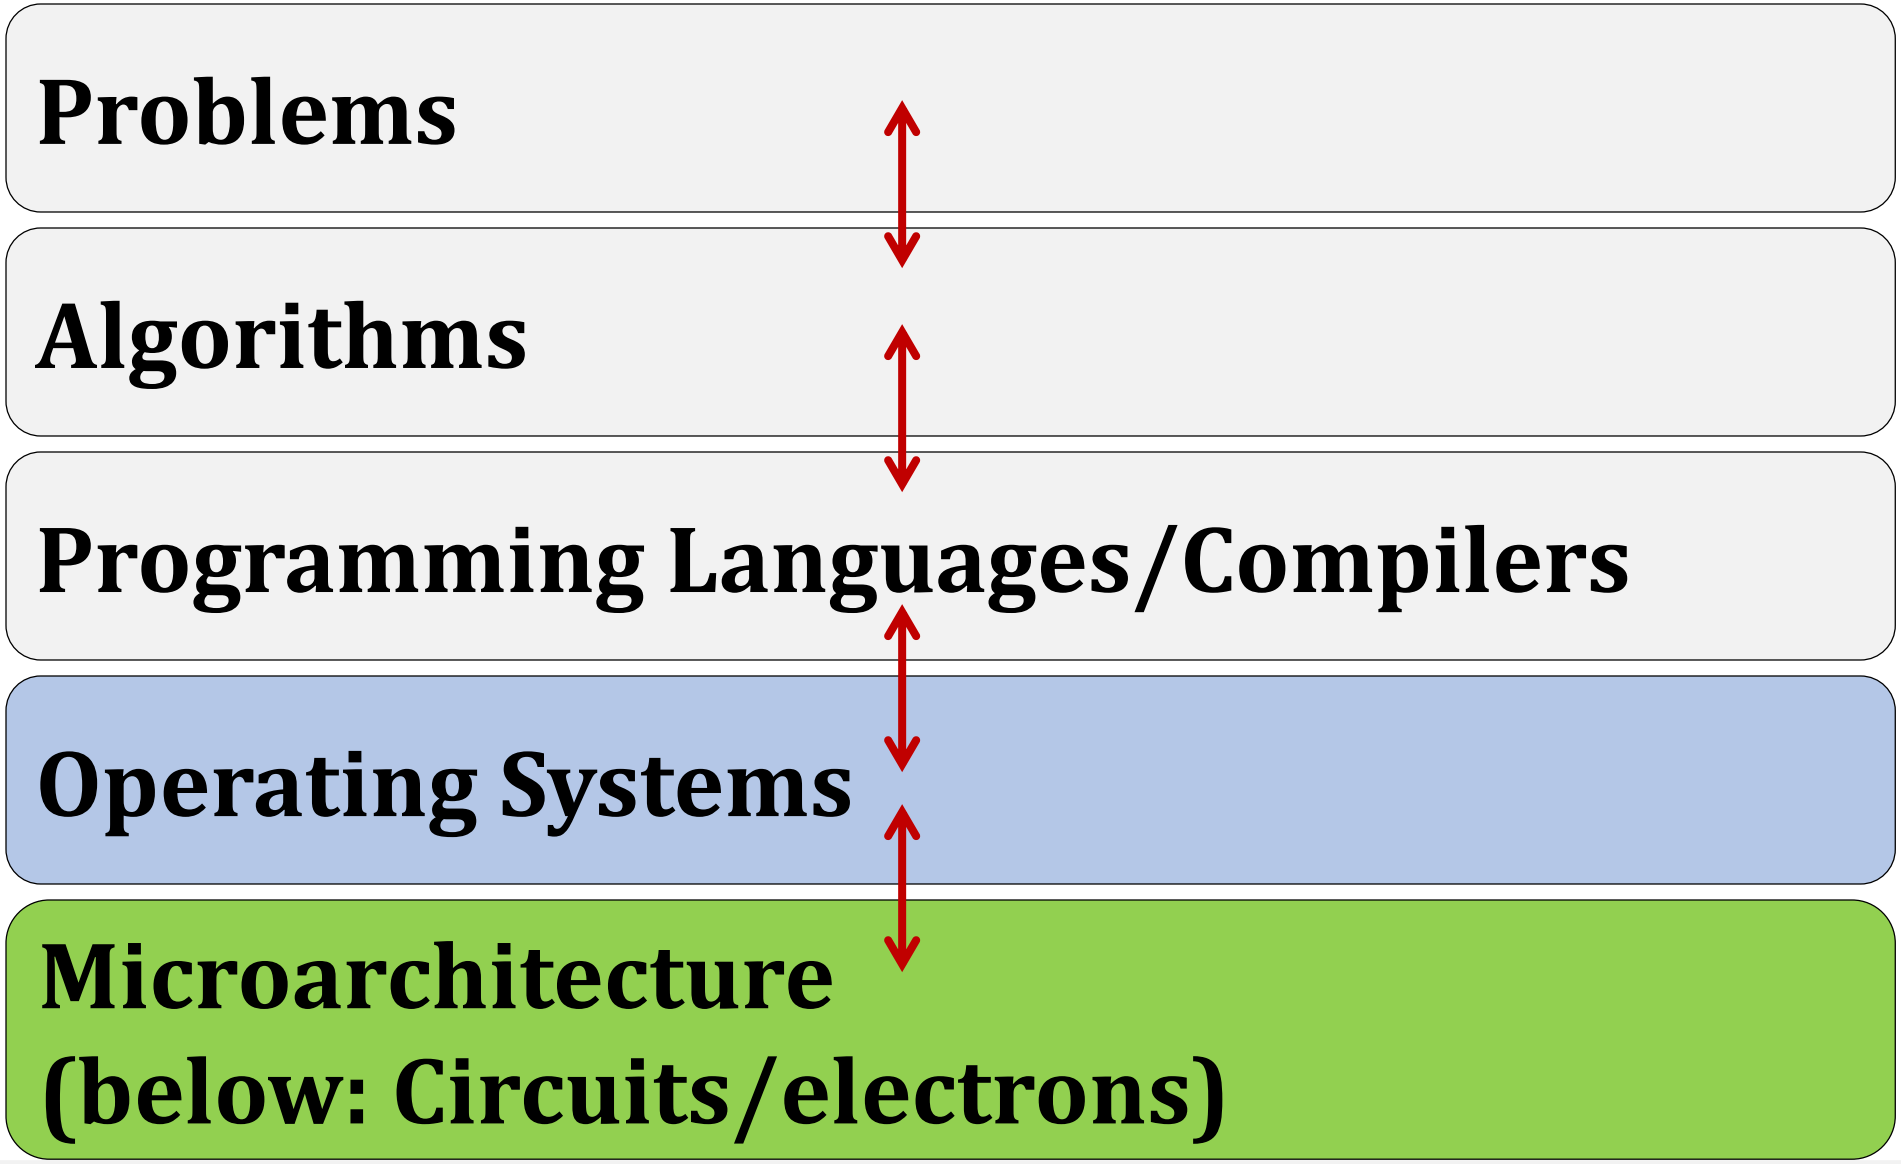

Mantra from Y. Patt [U.T. Austin]

Look Backward: Examine Old Code

Look Forward: New domains and new challenges

Look Up: Nature of Problems in the stack

Look Down: Technology (have EE friends)

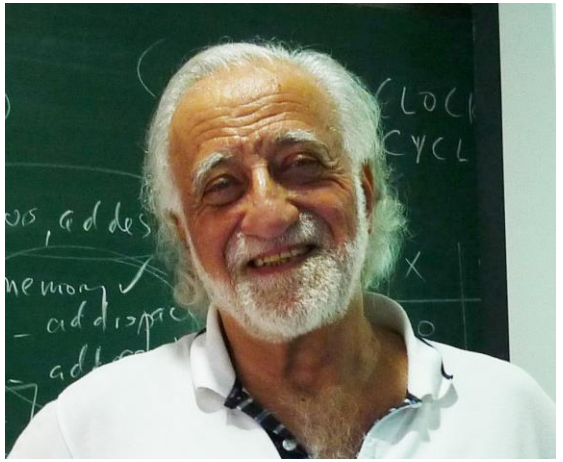

Let's Get Started Then

Assignment 0.0: Due midnight (tonight)
Link: Course Web-page
On-time submission: 0 point
Late submission: -1 point

Let's Get Started Then

Assignment 0.1: Due Jan 8th (1.00 P.M.)
What: Spend 20 mins on a video and answer few Qs.
Link: Course Web-page
On-time submission: 0 point
Late submission: -1 point

Brush-up ISA before next lecture: Appendix A in H&P

THANKS