


PLENTY OF ROOMS@COMPUTER SYSTEMS

Biswa@CSE-IITB

ACM ROCS 2024 Feb. 24, 2024

Who Said This?

"Why cannot we write the entire 24 volumes of the Encyclopedia Britannica on the head of a pin?"

CLUE?

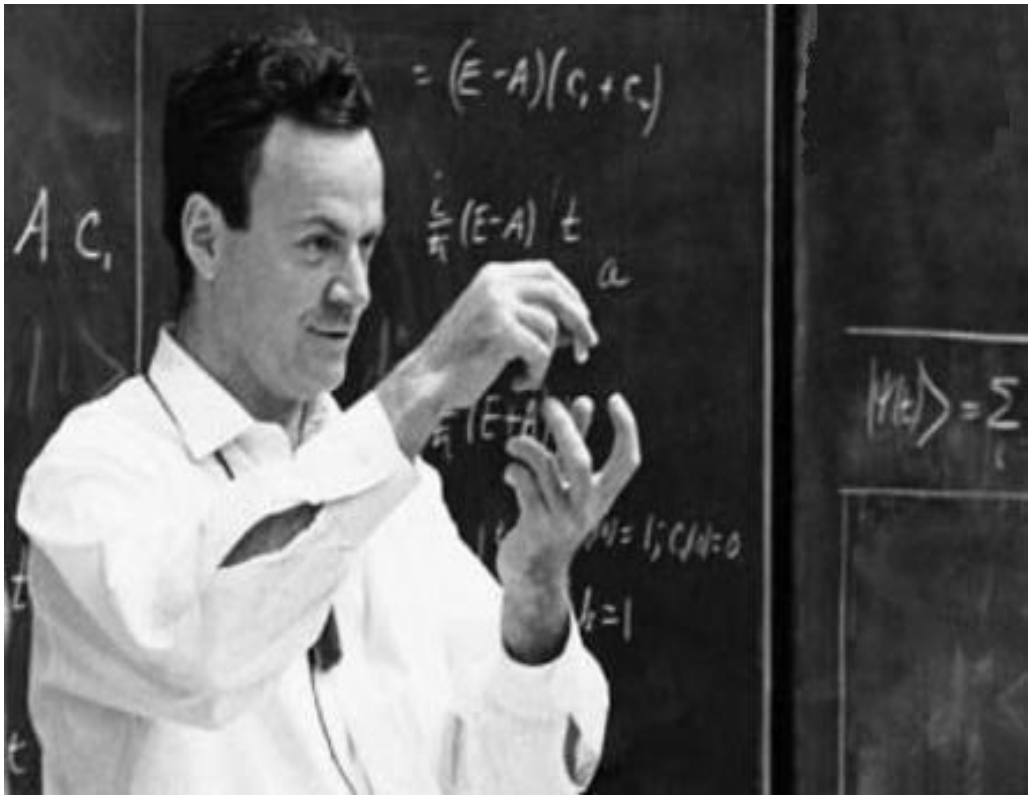

Bongo Player ??

Physicist ??

Noble Laureate ??

RICHARD FEYNMAN

Talk delivered in 1959

The question: 1000$ Challenge

Head of a pin $1/16^{th}$ of an inch.
Magnify it 25000 times and
that is sufficient area for all the pages in the encyclopedia.

Today's Talk – A Curtain Raiser

Room in Computer Systems: *Plenty* of it.

What is (in) Computer Systems ?

After the talk: Go and talk to people of interest to find out problems of interest.

What is Computer Systems

Systems starts where theory ends

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

CSE: Science (Theory) + Engineering (Systems)

Confused? Let's use high school .....

Theory: Value of "g" is 9.8 m/s^2

Systems: Value of "g" is 9.8 m/s^2 ?

Seeing/doing/experimenting is believing

The Key to Systems Research

Experiments, experiments, and experiments

Time for a video

https://www.youtube.com/watch?v=OL6-x0modwY

More on Experiments (Feynman again)

If a theory/idea does not agree with the experiment, then something (theory/idea) is wrong ☺

It does not make a difference how beautiful your idea is, how smart you are, how famous you are .....

WRONG

Welcome to the world@ Computer systems

The Computing Stack

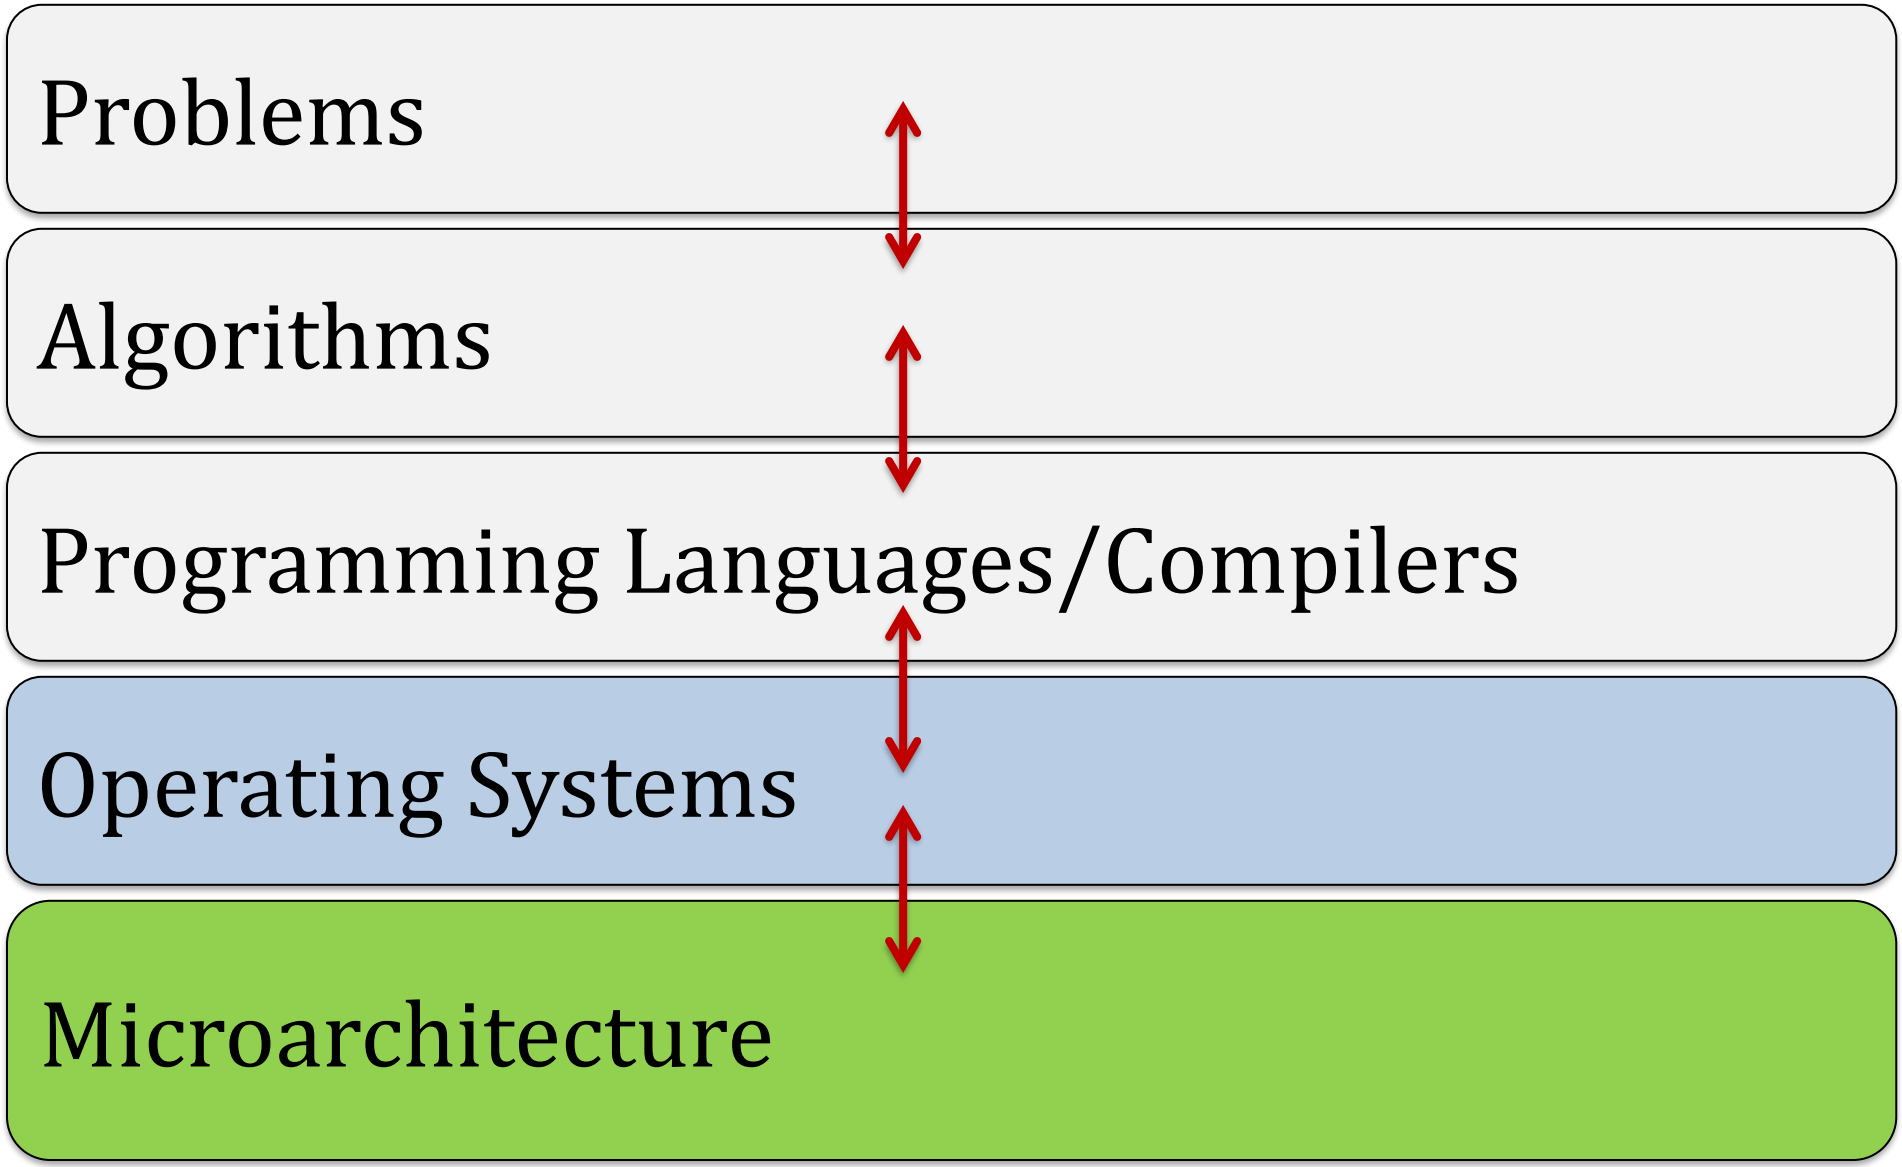

Computing to Communication/storage

Computer Networks

Distributed Systems

Cloud based Systems

Database systems

Computer Architecture

Compilers

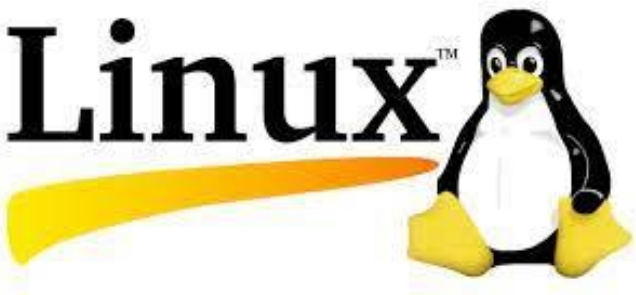

Operating Systems

Computer Networks

Hey, this is a workshop on research?

Research: More and more on less/less

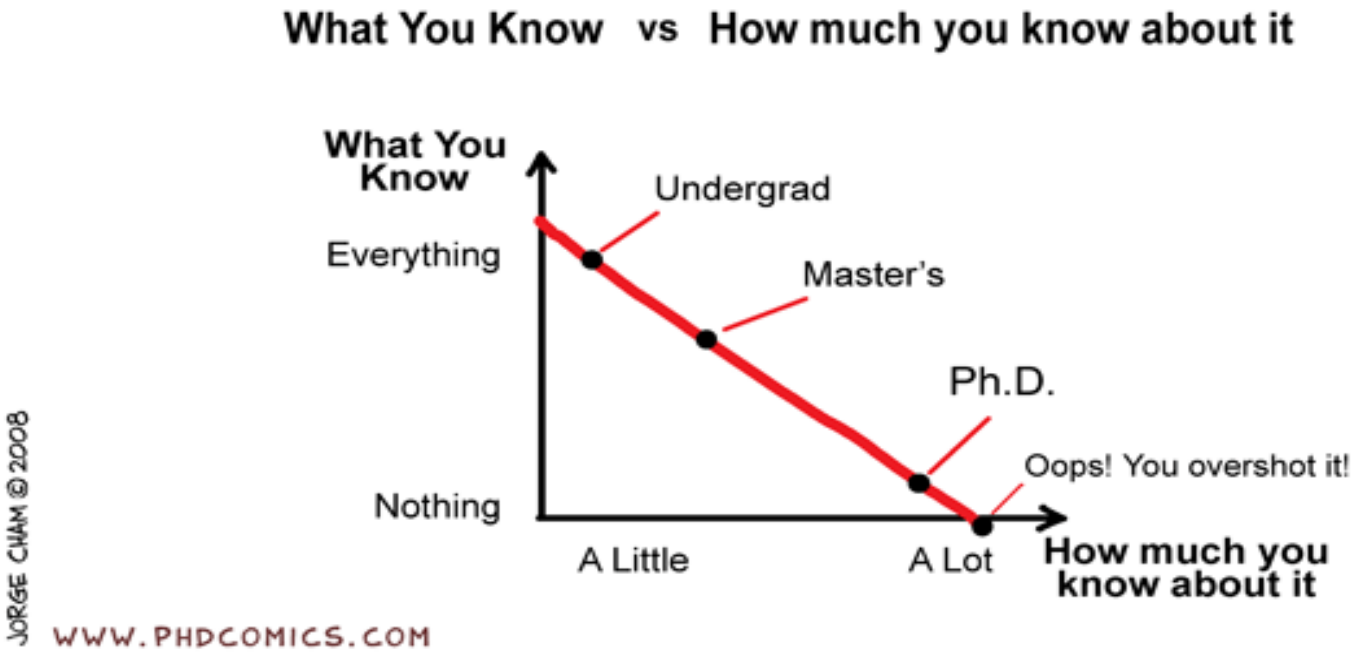

Human knowledge

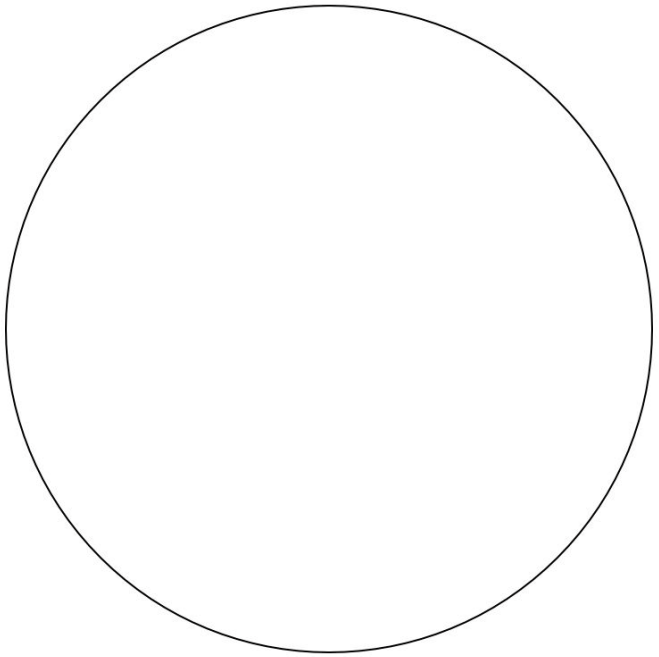

After school

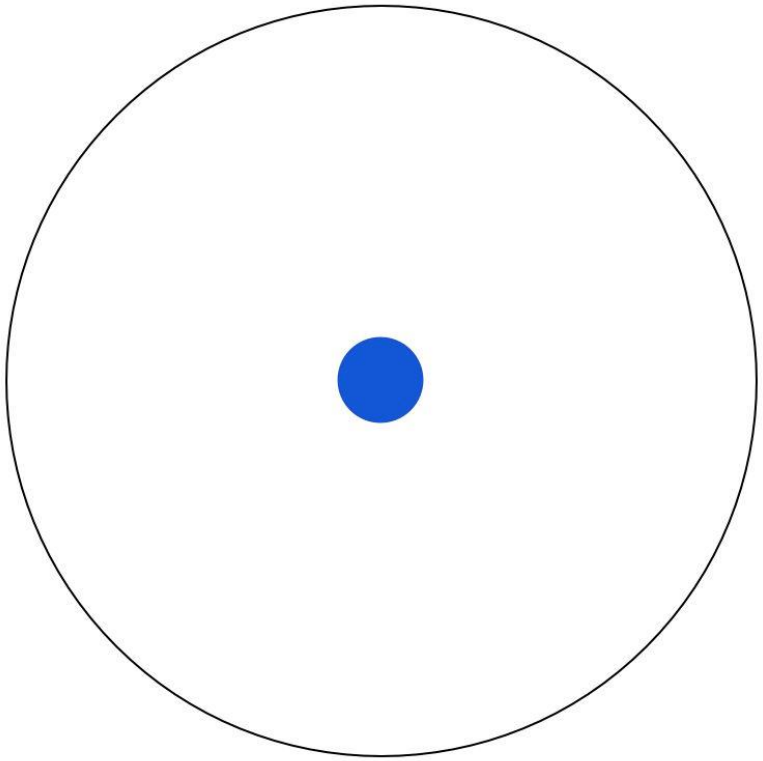

After B.tech.

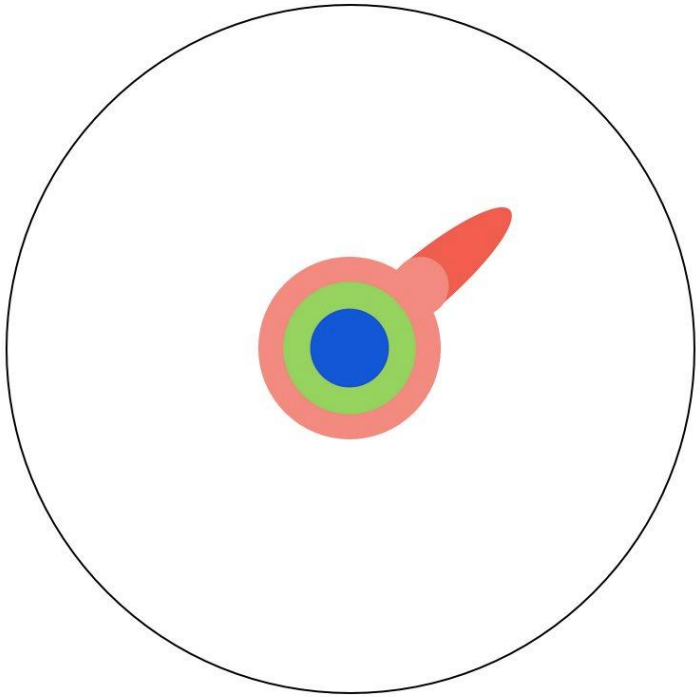

After interest in computer systems

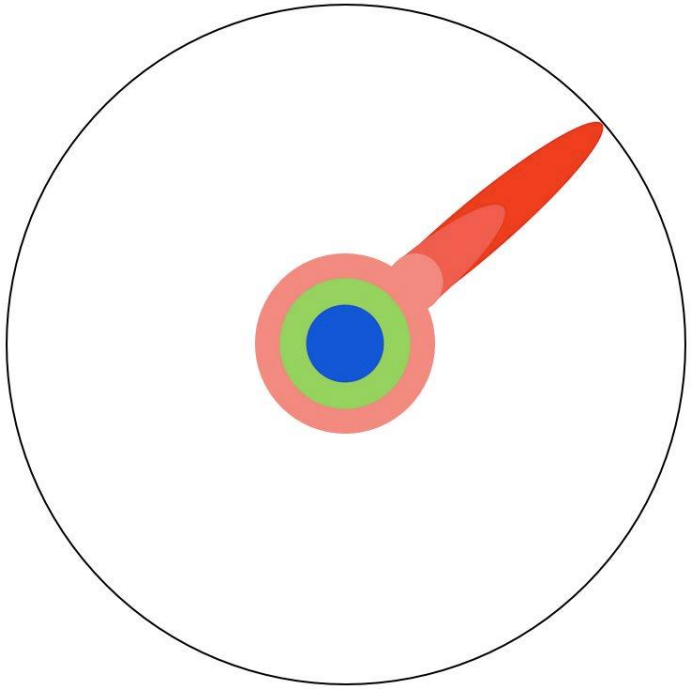

One fine day

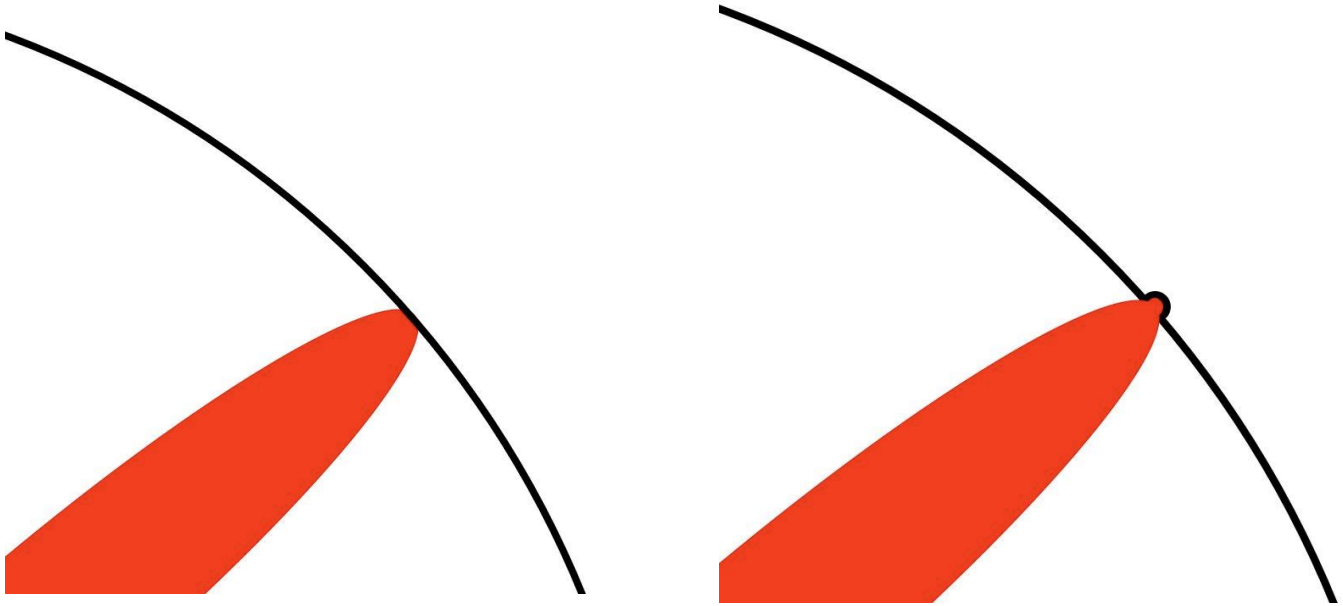

The bigger picture

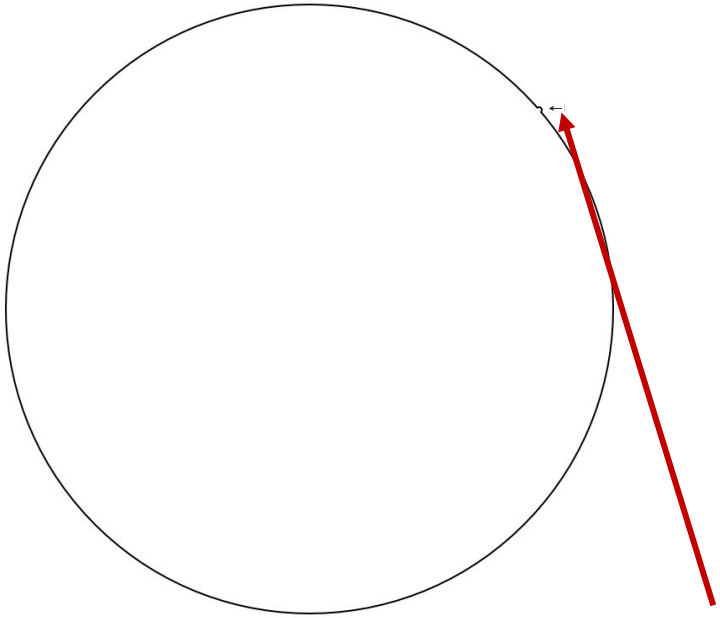

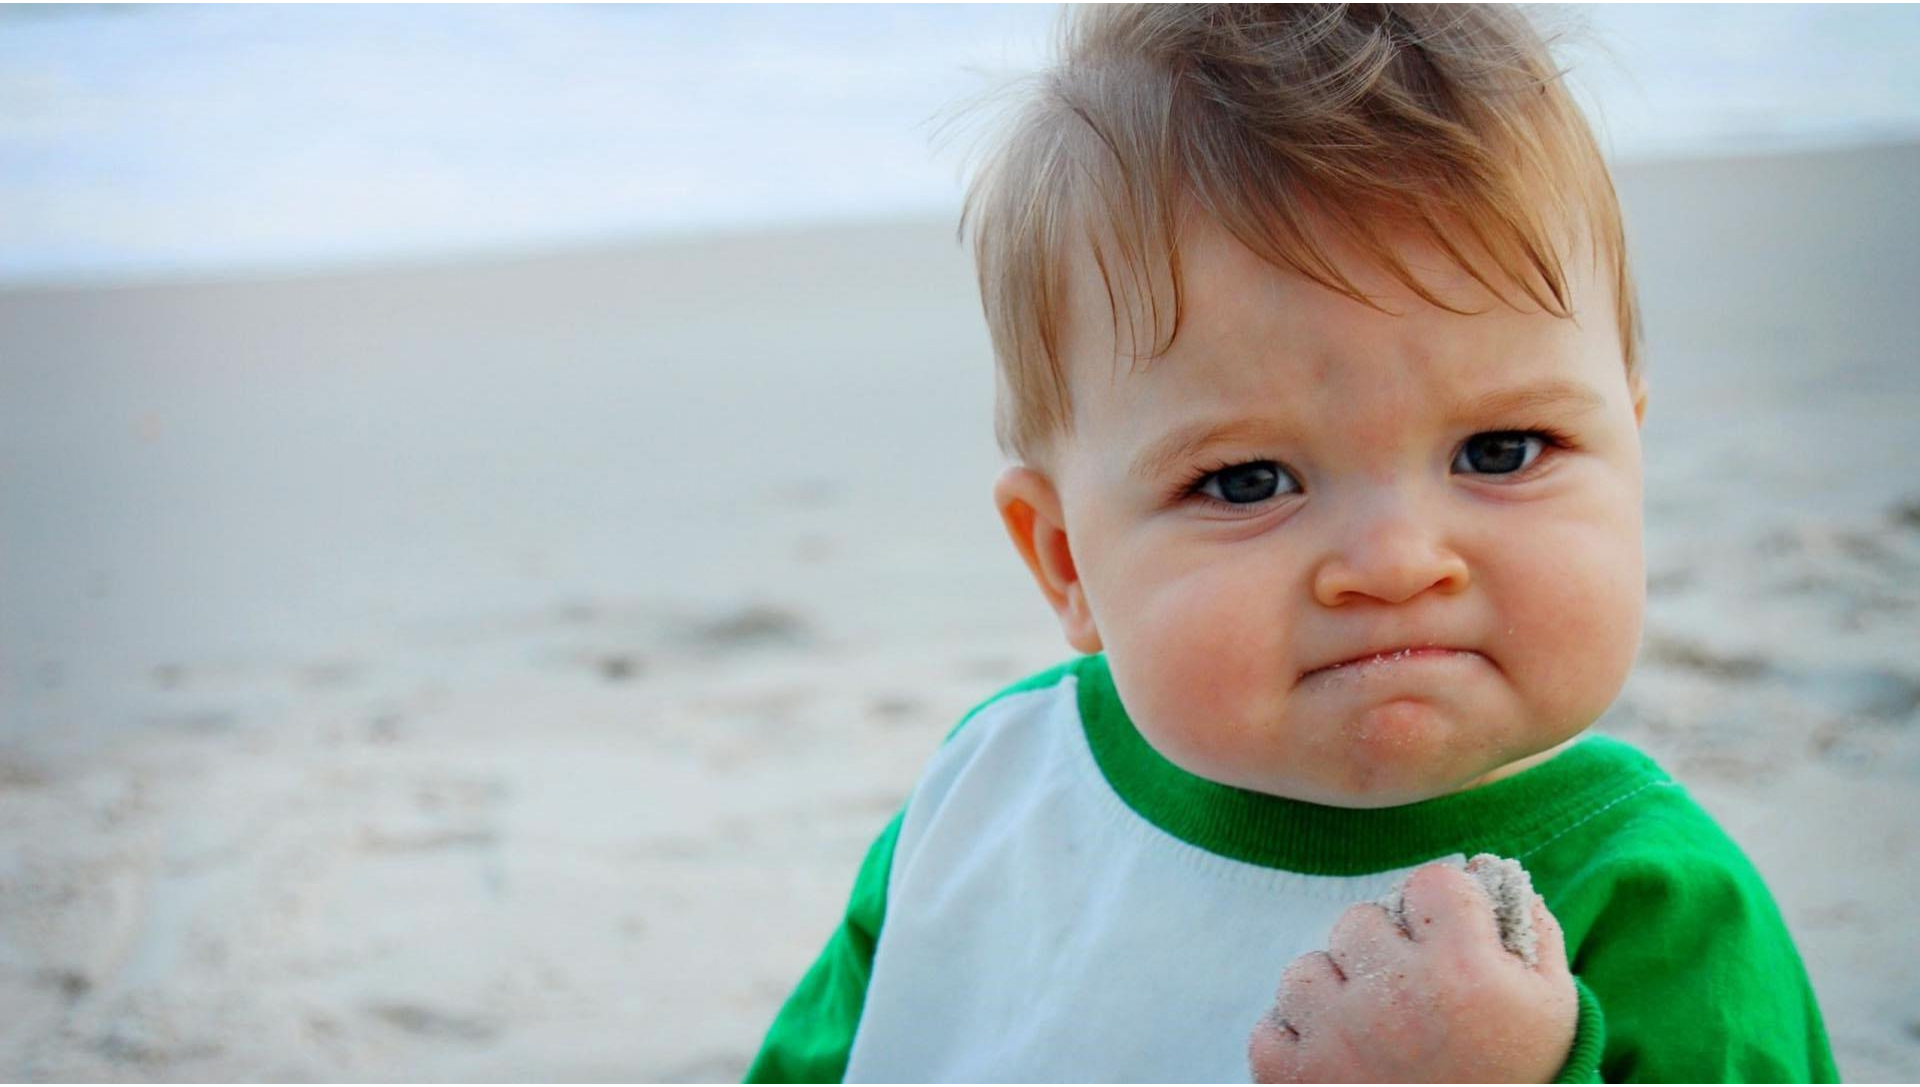

The joy

Idea/implementation named after YOU

You are a celebrity overnight

You are the producer of knowledge and not the consumer

Computer Systems Research

Hum sath sath hain

PL+compilers+architecture+
OS+networks+database

Cohesive ecosystem that makes our life easy

Computer Systems Tree

Why Computer systems research?

Enabler of all areas

It is exciting and it is everywhere

Think about AI/ML,
New age computing helped ☺

Computer Systems Research

Tradeoffs

Metrics of interest: Sweetspot

Performance

Power

Cost

Complexity, scalability, flexibility, etc

10K view of a systems researcher

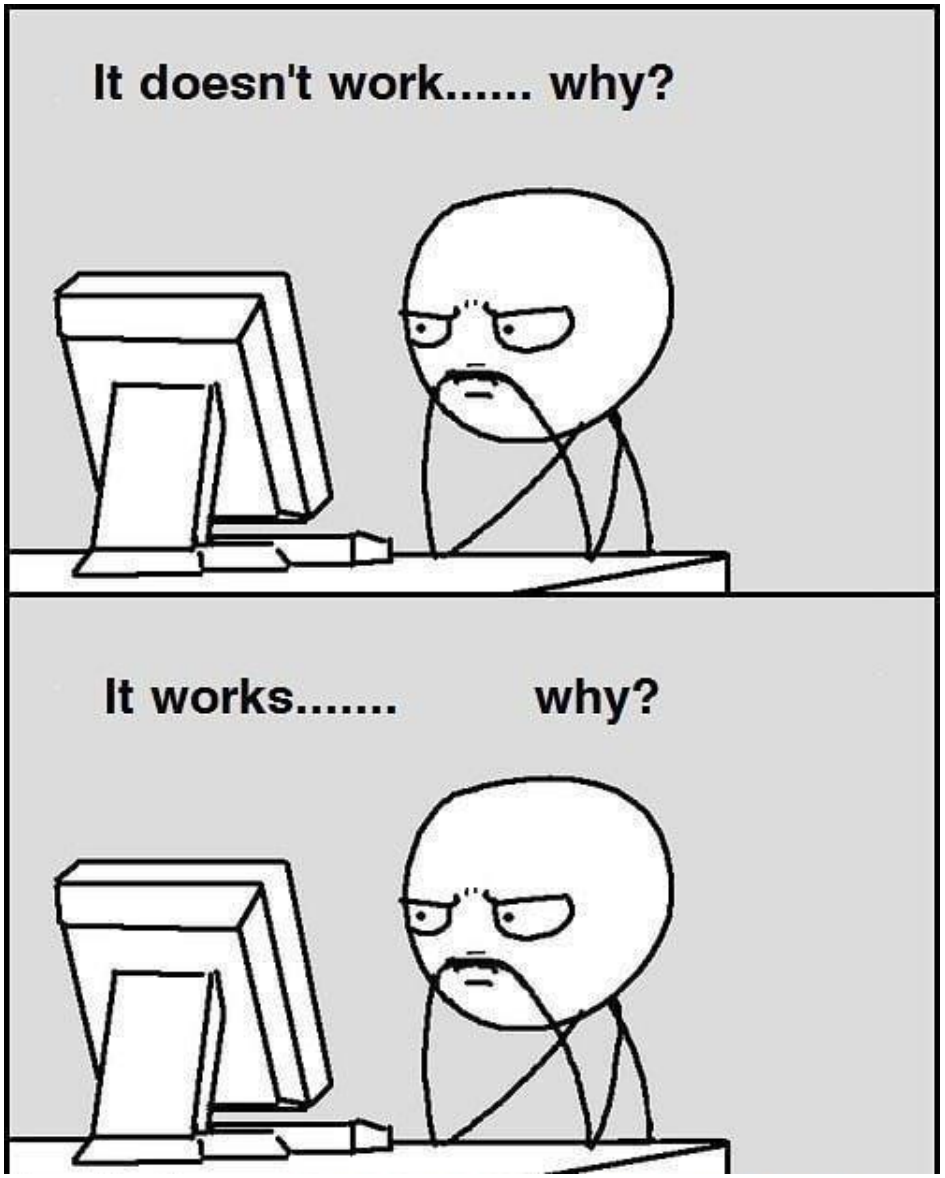

Works, does not work?

Can I design a processor that gets all the data from memory in nanoseconds?

Can I design an OS scheduler that can work seamlessly for a 100-core system?

Can I come up with optimizations that will reduce the size of a binary?

How to design memory systems for 100Gbps network I/O?

AI/ML vs Systems Research (Effort)

My Systems Papers (N > 1) vs My NeurIPS paper (N =1)

Example: Usenix ATC-2020, Usenix Sec-2020

(For each paper)

- 3 students, total 3,000 man hours of implementation and experimentation
- 3 math formulae, all understandable by a high school graduate
- 10+ result figures
- 5000+ lines of code (LOC)
- 1.5 years from idea to paper accept
- Had to understand design and replicate 3 prior papers
- 16-18 pages
- One of 65 and 157 accepted papers

NeurIPS-2020

- 1 student, total 200 man hours of implementation and experimentation
- One super cool intuition
- Three formulae, none of which is understandable by a high school graduate, one I barely understand and only because my ML faculty co-author explained it at length
- 2 result plots, 1 result table
- 200 LOC
- 6 months from idea to paper accept
- Had to understand math and replicate 3 prior papers
- 10 pages
- One of 1428 accepted papers

Prof. Saurabh Bagchi, Purdue university

Mantra for Systems Research

Abstraction? Break it

It is good if you don't care about the performance of underlying entities.

What?

ABSTRACTION BARRIER

How? Why?

How many of you can drive a bike?

How many of you know how a bike works?

Systems: How bike works?

So break your abstraction barriers

Can you break the abstraction barriers?

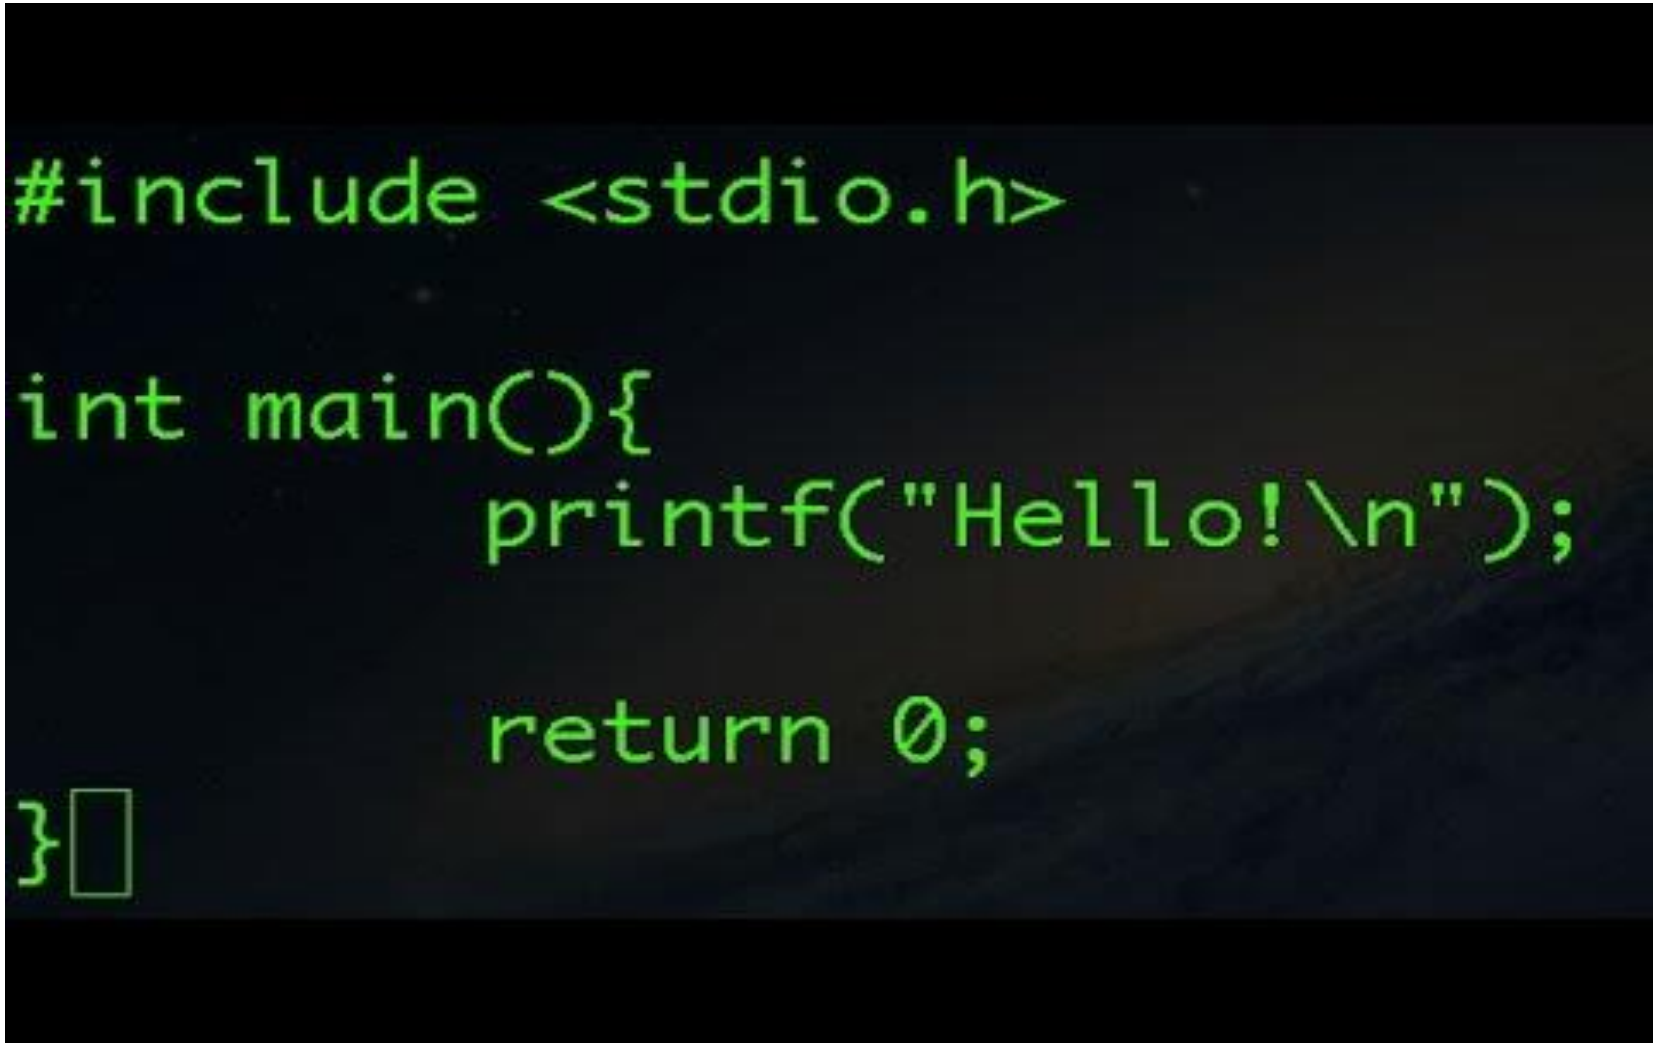

Computer Systems: 10,000 feet view

Program written in C, compiled using gcc

Scheduled by OS, executed by processor, by fetching data from memory

Communicated to systems (may be located far away) through network protocols

Ensuring secure execution and communication, Verifying whether systems work the way it is supposed to work (robustness)

How Can You Help?

Understand the layers/barriers of abstraction

Know the HOW and question the WHY ?

Can you make it better?
Think Big: *by thinking about small things*

Build systems for future

Theorems to Tools

Hack real OS, compiler

Play with real and simulated hardware

Errors, bugs, ...

Make a pitch quantitatively (experiments)

Top conferences: Look for csrankings

Ranks top conferences only

Systems [off | on]

- ▶ Computer architecture ☑
- ▶ Computer networks ☑
- ▶ Computer security ☑
- ▶ Databases ☑
- ▶ Design automation ☑
- ▶ Embedded & real-time systems ☑
- ▶ High-performance computing ☑
- ▶ Mobile computing ☑
- ▶ Measurement & perf. analysis ☑
- ▶ Operating systems ☑
- ▶ Programming languages ☑
- ▶ Software engineering ☑

Reach out to Profs for research internships (six to eight months)

We need the next gen. who can push the limits .....

https://csrankings.org/#/index?all&in

One Step at a Time

Go through top quality course lectures online

Go through webpages of profs of interest

Online talks: Systems@India talk series: check my webpage

Join for research internships or masters/Ph.D. programs

What you need? 3S (My Take)

Shradha

Sadhana

Sahas

Once you are in 3S zone: Know the problem to know the solution

Key Takeaways

Understand/break/re(develop) the abstraction barriers

Experiments, experiments, and experiments

If something works, why? && 3S

Famous Quotes ☺

*"I don't think there will ever be a market for more than **five** computers in this world"*

*"**640 KB** of memory ought to be enough for anybody"*

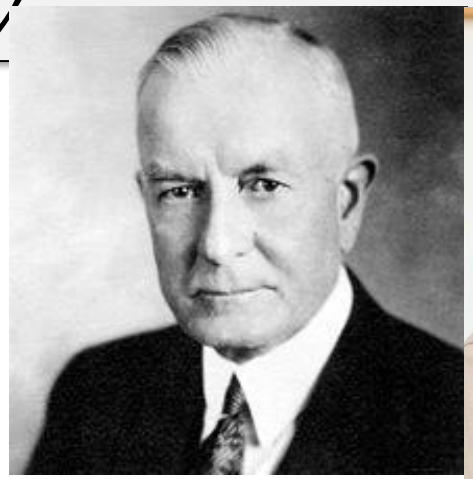

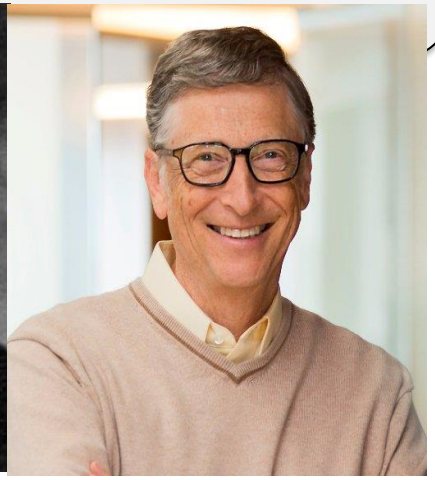

Reality !!

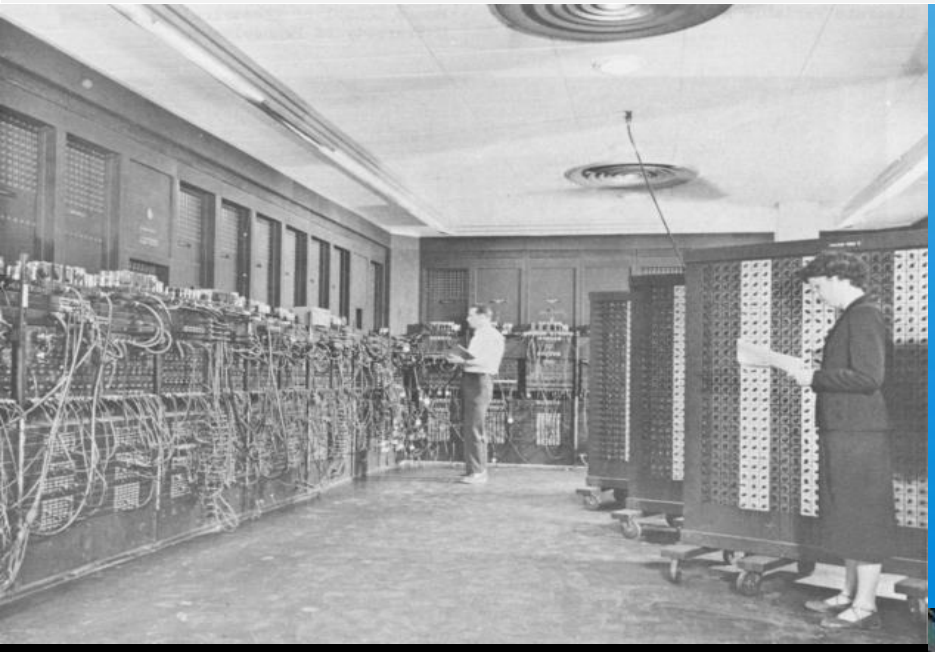

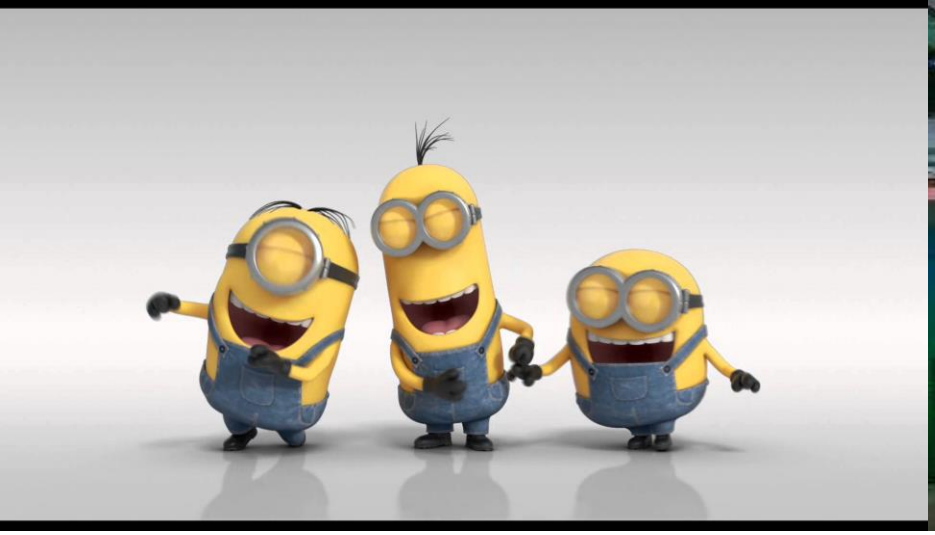

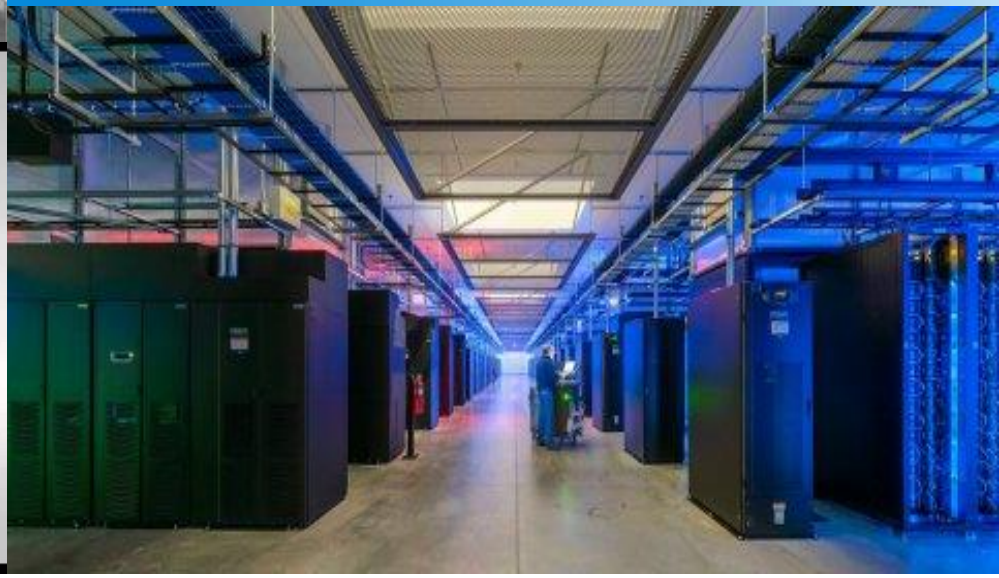

"It takes two to speak the truth - one to speak and another to hear" - Henry David Thoreau

Thank You

Reach me: biswa@cse.iitb.ac.in

More info. On systems research

https://www.cse.iitb.ac.in/~biswa/

CASPER group

https://casper-iitb.github.io/