The Awesome World Enabled by the Transistor

Alvin Loke

NXP Semiconductors
San Diego, CA

Special Acknowledgment to
Alessandro Piovaccari

Semiconductors in the News

Must be globally very important

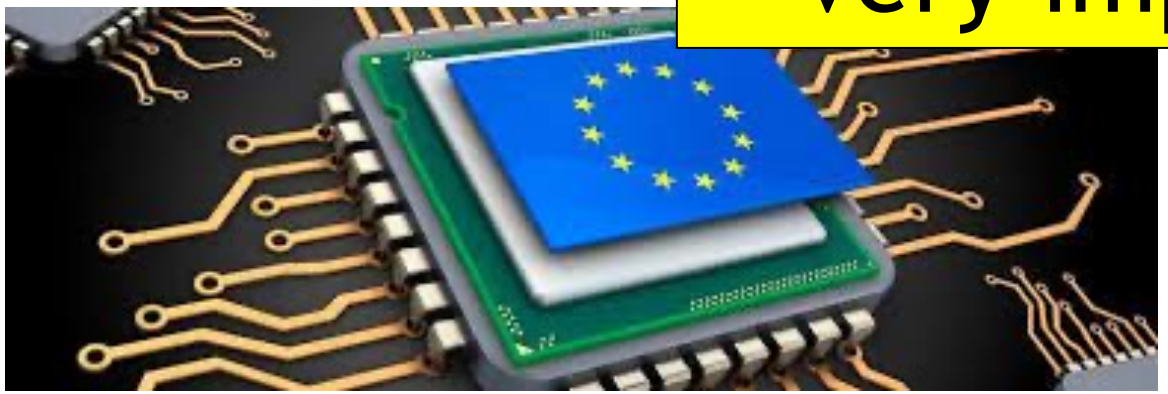

A Humble Beginning 75 Years Ago

75TH ANNIVERSARY | TRANSISTOR

- First transistor was successfully demonstrated on December 23, 1947 at Bell Labs in Murray Hill, NJ (research arm of AT&T)
- Invented by John Bardeen, Walter Brattain & William Shockley

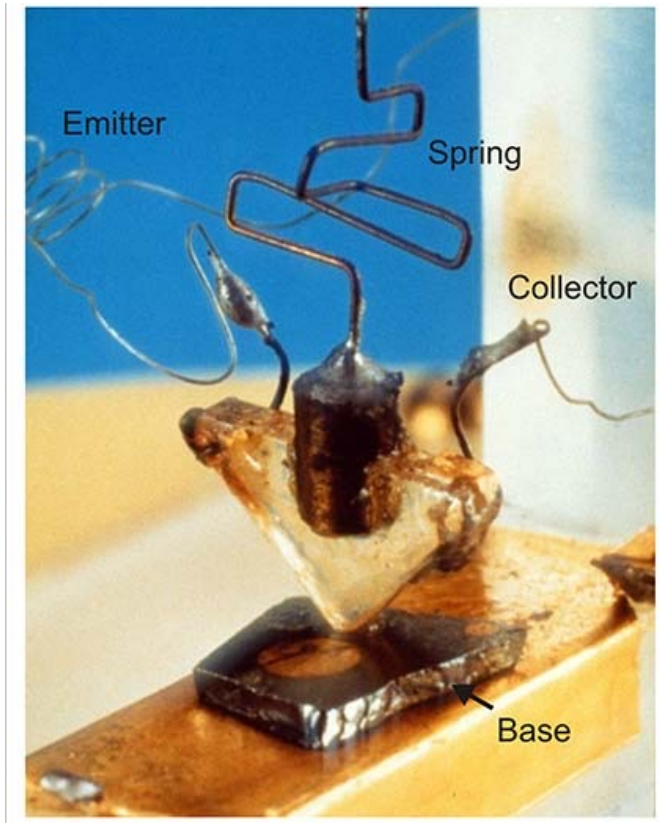

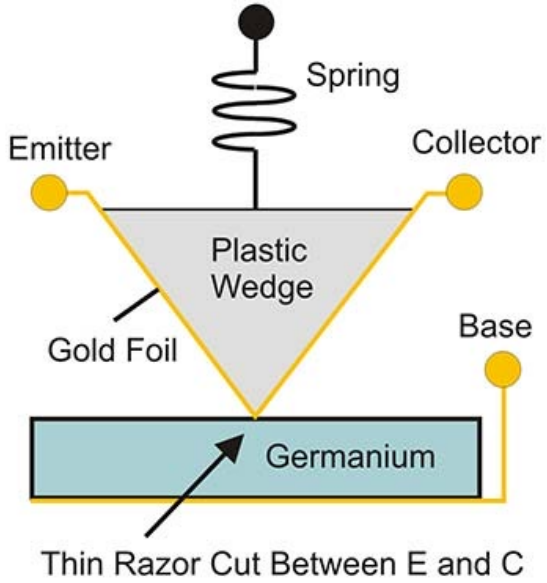

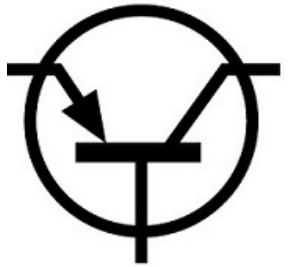

https://www.nutsvolts.com/magazine/article/the-story-of-the-transistor

Me In a Nutshell

- Born in Malaysia, grew up in Vancouver (Canada), adult & family life in US
- Education: UBC BASc Engineering Physics, Stanford MS/PhD Electrical Eng.
  - Tinkered with electronics since 8th Grade
  - Intern for 6 summers (Texas Instruments, Motorola, Sumitomo Electric, …)
  - Fell in love with semiconductor physics in college junior year, still in love with it
- 24 years in industry (HP/Agilent, AMD, Qualcomm, TSMC, NXP)
  - Currently NXP Technical Fellow
  - Started in semiconductor technology development
  - Moved to analog design & technology/modeling interface
  - Now focused on design methodology & high-speed design
  - Worked on every CMOS node from 250nm down to 2nm (15 nodes)
  - Lived in Bay Area, Osaka, Singapore, Texas, Colorado, now San Diego
  - Active IEEE volunteer for 20 years
- Wife Tin Tin is also circuit designer with semiconductor technology background
- Two kids – Theo (9th Grade – go Ravens!) & Josephene (6th Grade)

My Mentors and Teachers


Matt Angyal


Larry Bair

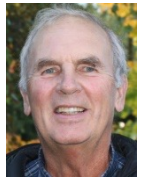
Bob Bardsley

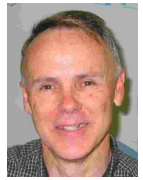
Bob Barnes


Suzanne Biganzoli


John Bravman

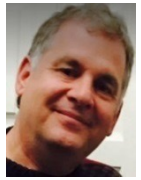
Mike Brunolli


Min Cao


Luigi Colombo


Tom Cynkar


Reinhold Dauskardt


Dick Dowell


Bruce Doyle


Pat Drennan


Emerson Fang


John Faricelli


Dennis Fischette


Dan Friedman


Doug Garrity


Lixin Ge


Mike Gilsdorf


Jung-Suk Goo


Bob Havemann


Rick Hernandez


Paul Holdaway


Tim Hollis


Mark Horowitz


Reza Jalilizeinali


Philippe Jansen


Ted Kagetsu


Ron Kennedy


Takamaro Kikkawa


Xiaohua Kong


Greg Kovacs


Phil Fisher


Joachim Kruecken


Steve Kuehne


Benny Lai


Ralph Lang


Mike Leary

C.K. Lee

Tom Lee

Justin Leung

Tom Lii

Rich Liu

John Long

Linus Lu

Shawming Ma

Colin McAndrew

Charles Moore

Don Morris

Bich-Yen Nguyen

Ken O

Michael Oshima

Chintamani Palsule

Marcel Pelgrom

Jim Pfiester

Alessandro Piovaccari

Jim Plummer

Dave Pulfrey

Gary Ray

Behzad Razavi

Jeff Rearick

Ken Rim

Changsup Ryu

Krishna Saraswat

Shawn Searles

Ray Stephany

David Sunderland

Gerry Talbot

Tom Tiedje

Paul Townsend

Jan van der Spiegel

Ram Venkatraman

Martin Wedepohl

Tin Tin Wee

Andy Wei

Jeff Wetzel

Daniel Weyl

Philip Wong

Simon Wong

Bruce Wooley

Joanne Wu

Philippe Wyns

Guoqiang Xing

Joseph Yan

Geoffrey Yeap

Bo Yu

Patrick Yue

Carl-Mikael Zetterling

Semiconductor Demand

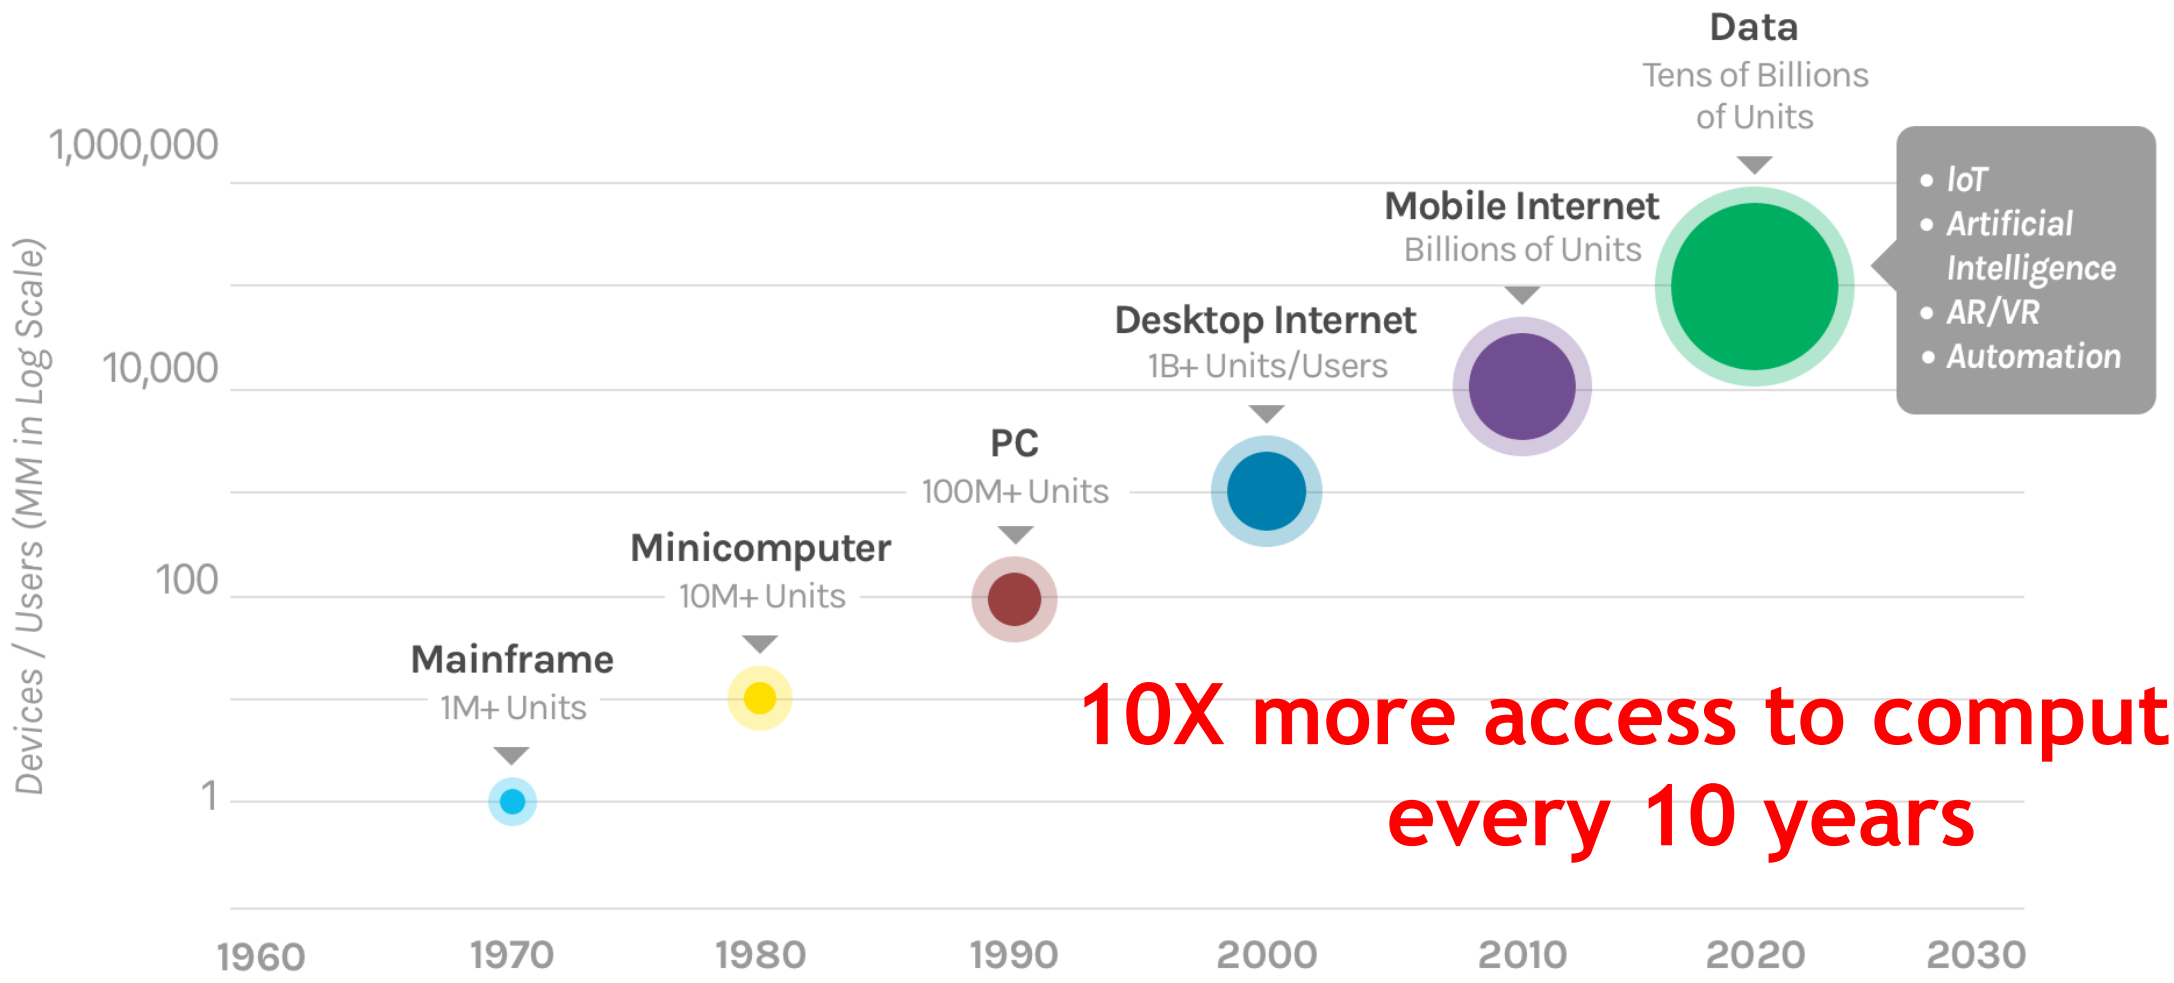

10X more access to computing every 10 years

Source: Morgan Stanley
(Courtesy Alessandro Piovaccari)

Worldwide Semiconductor Market Size

Let's put it in context…

→ Worldwide Markets (2021)

- Semiconductors: $553B
- GDP: $93,864B
- IT Data Centers: $228B
- IT Devices: $705B
- Car & Auto: $3,600B
- Home Appliances: $420B
- Semiconductor market expected to hit $1T in 2030

Sources: Statista, IBIS World, SIA
Reference: [SIA2021], [Gartner2021]
(Courtesy Alessandro Piovacarri)

What Makes This All Possible?

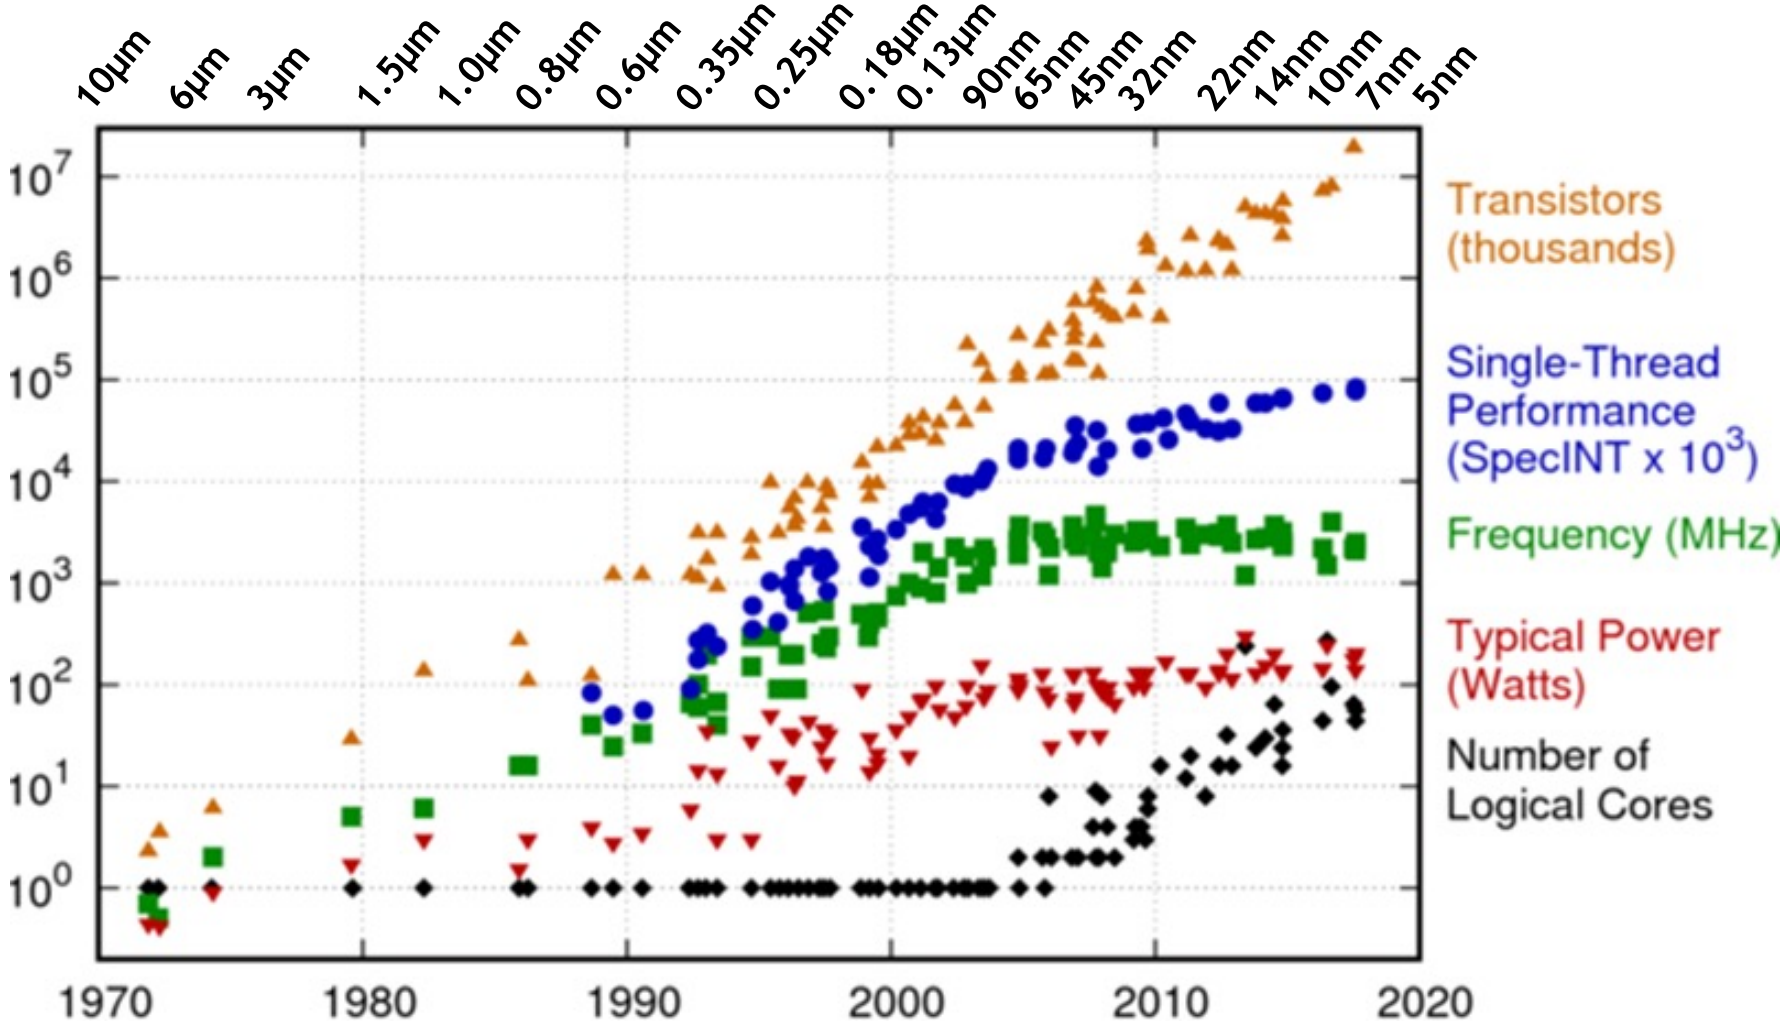

Source: https://www.karlrupp.net/2018/02/42-years-of-microprocessor-trend-data/

Miniaturizing (scaling) transistors makes them cheaper, faster, and more energy efficient

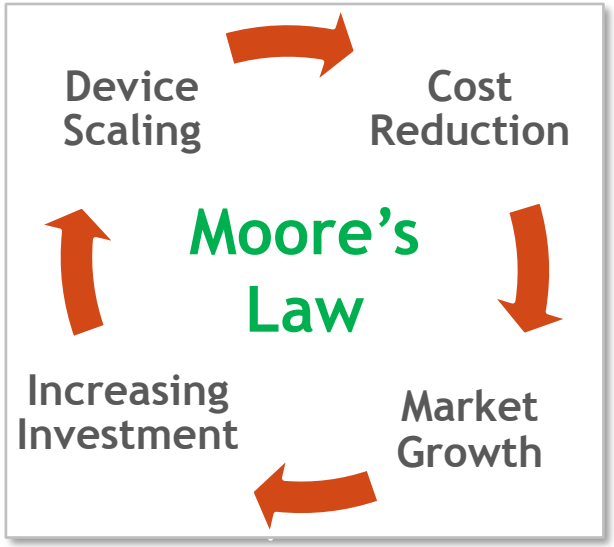

(Courtesy Alessandro Piovacarri)

What 50 Years of Moore's Law Has Enabled

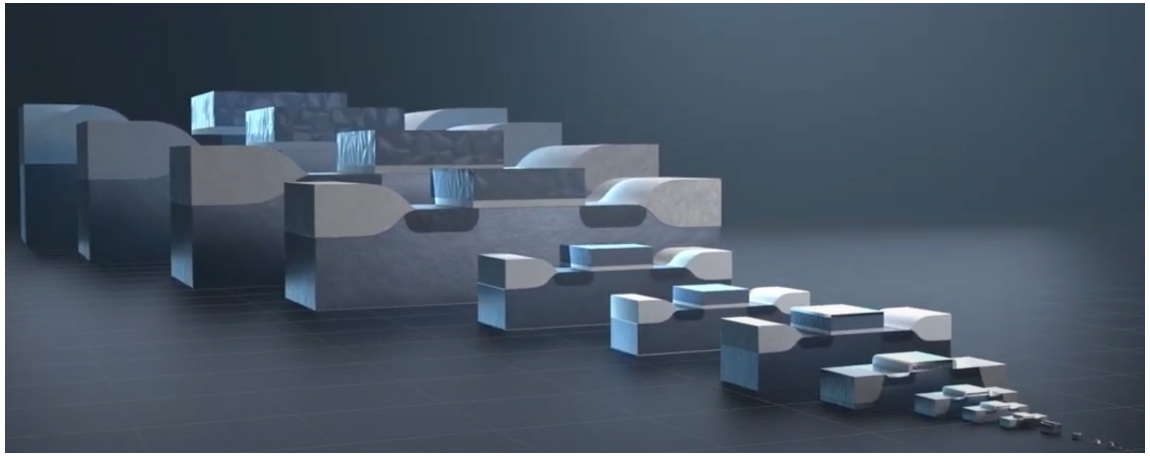

(Source: intel.com)

- 1,000,000X smaller
- 3,500X greater performance
- 60,000X lower cost
- 90,000X more energy efficient

If car technology progressed at the same pace as semiconductors, the VW Beetle would:

- Go 300,000 mph
- Cost $0.04
- Get 2,000,000 miles per gallon of gas
- Last your entire life on one single tank of gas

(Courtesy Joe DiFilippo)

What's Possible After 50 Years

AMD Xiaolong (2021)
32B transistors (38.5 total)
TSMC 7nm + GF 14 nm + organic
64 cores - 256MB cache
2.25-3.4 GHz - 165-225 W

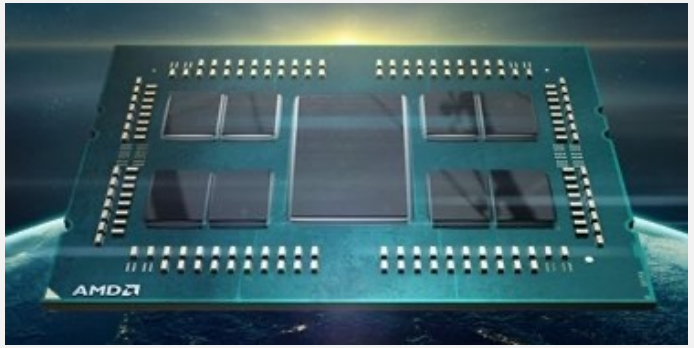

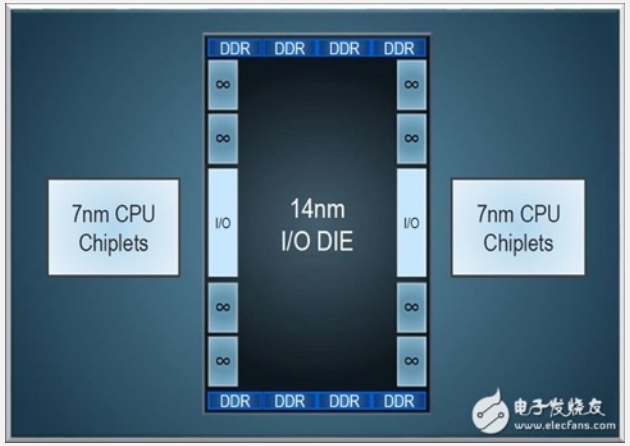

Source: *firstxw.com*

Xilinx Virtex Ultrascale+ VU19P (2019)
35B transistors (total)
TSMC 16FF+ + CoWoS (4 dies)
16xA9 - 9M SLCs
2K I/Os - 4.5 Tbps BW

Source: *AnandTech*

Apple M1 Max (2020)
16B transistors (57B with RAM)
TSMC 5nm (core)
8 CPU + 8 GPU - 16GB RAM
3.2 GHz - 7-39 W

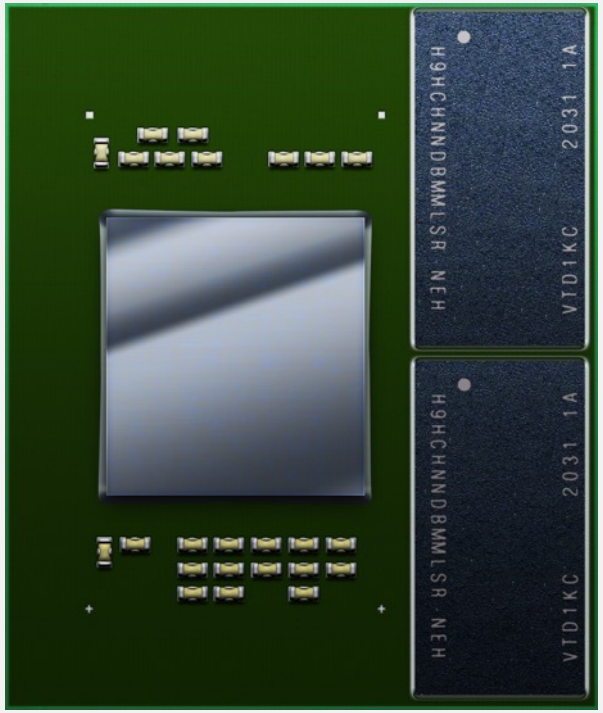

Source: *wikipedia.com*

(Courtesy Alessandro Piovaccari)

More Amazing Possibilities

Tesla D1 (2021)
50B transistors (total)
TSMC 7nm
425MB cache - 16 Tbps BW
362 Tflop (BF16/CFP8) - 400W

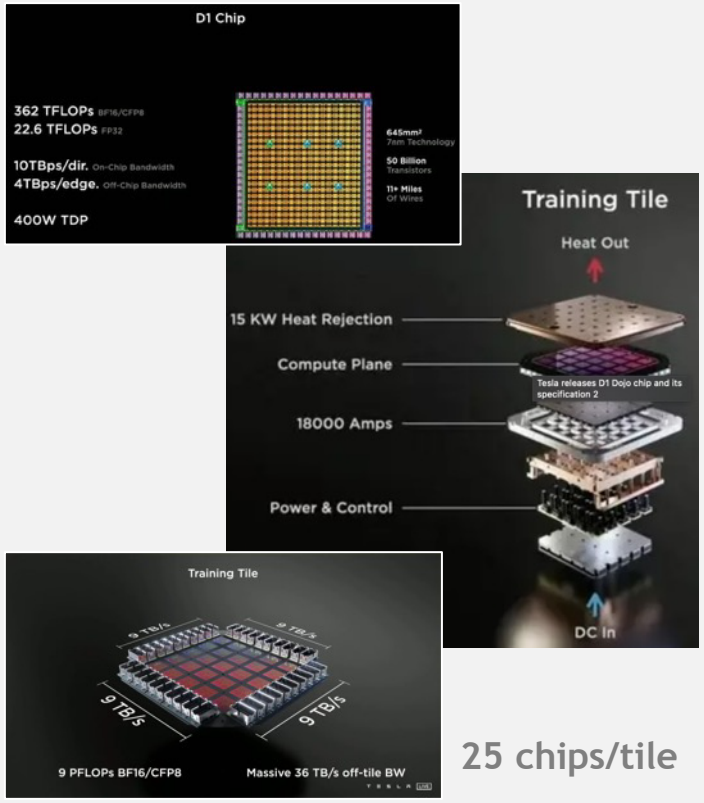

25 chips/tile

Source: teslanorth.com

(Courtesy Alessandro Piovaccari)

Samsung V-NAND Flash (2020)
2T transistors (3D stack)
Samsung V-NAND (100+ layers)
8x 256 Gb dies
1.4 Gbps

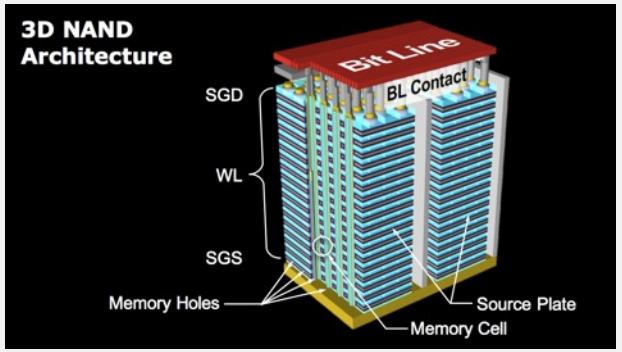

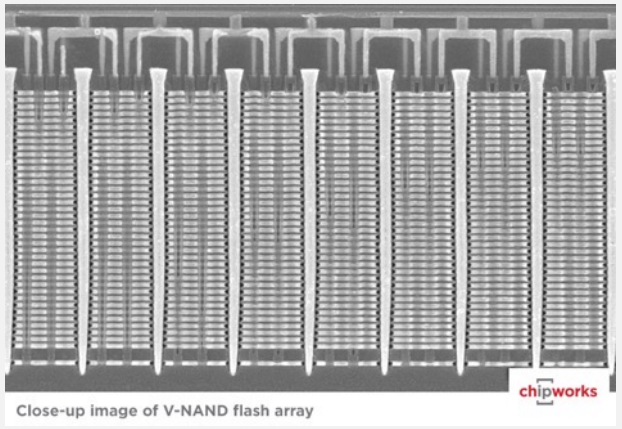

Source: Samsung, Chipworks

Cerebras WSE-2 (2019)
2.6T transistors
TSMC 7nm
850K cores - 40GB RAM
-

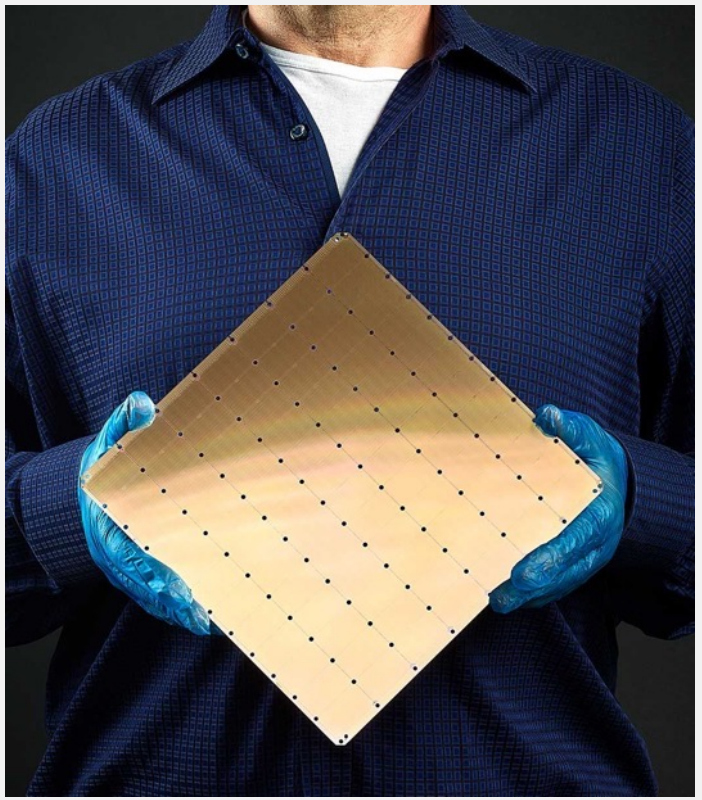

References: [Moore2021ISM]

A Semiconductor Fab is Like a Book-Printing Plant That Happens To Be WAY More Expensive

Human Scale

Book Printing	Chip Fabrication
An author writes a book *They use a word processor*	 An engineer designs a chip *They use EDA Tools*
They contract with a publisher who sends text to the printing plant *It may print novels, tech manuals, histories, etc.*	 They select a Fab appropriate for their type of Chip *Memory, logic, RF, analog*
The plant buys raw materials *Paper, ink*	 The fab buys raw materials *Silicon, chemicals, gases*
The plant buys printing machinery *printing presses, binding, trimming*	 The fab buys wafer fab equipment *Etchers, deposition, lithography, testers, packaging*
The printing process - offset lithography *Filming, stripping, blueprints, plate making, printing, binding, trim*	 Chip manufacturing process - offset lithography *Etching, diffusion, lithography, assembly, testing, packaging*
The plant turns out millions of copies	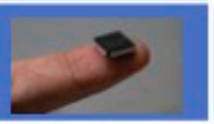 The plant turns out millions of copies

Atomic Scale

Source: SemiWiki

(Courtesy Alessandro Piovaccari)

A Modern Mega-Fab (Wafer Fabrication Plant)

- Typical price tag of $20B
- Throughput of 30-50k wafers per month
- Depreciates at $0.5M per hour

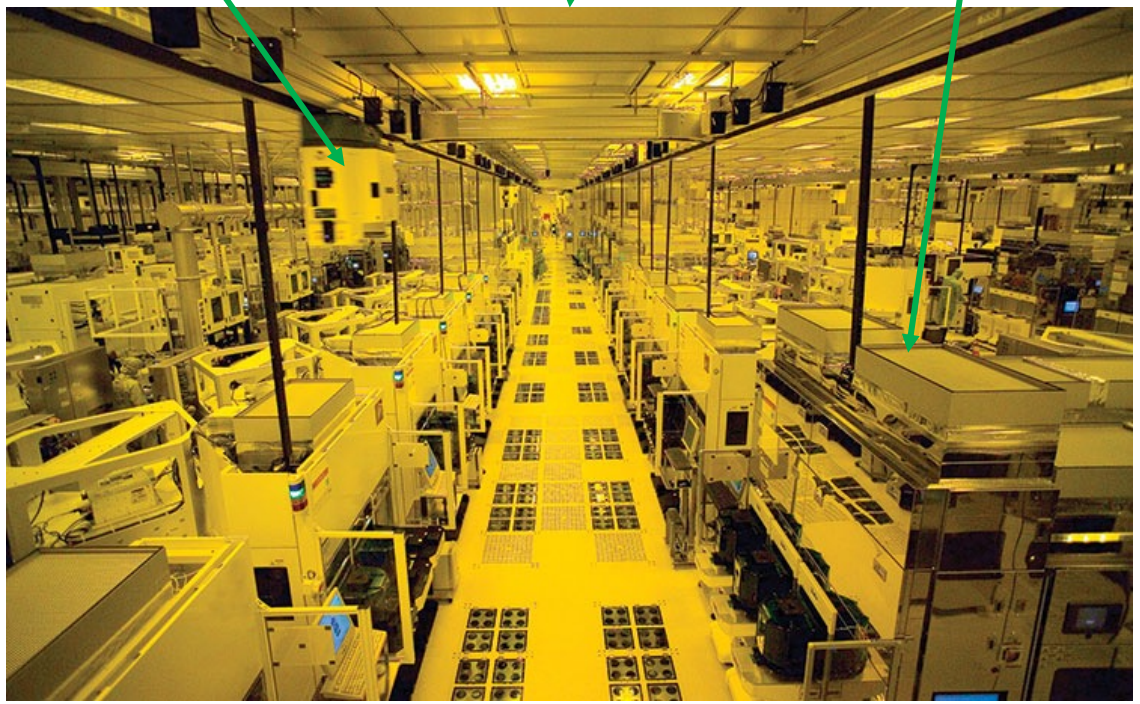

Source: WikiChip

(Courtesy Alessandro Piovaccari)

Chip Fabrication (Wafer Processing)

- All the action happens on the surface of a silicon wafer
- Sequentially deposit or remove/etch layer after layer of different materials
- Each deposited or etched layer is patterned using a different mask pattern
- Several thousand steps to complete a chip using 40 to 100 masks
- Chips are probably the most complicated products ever manufactured

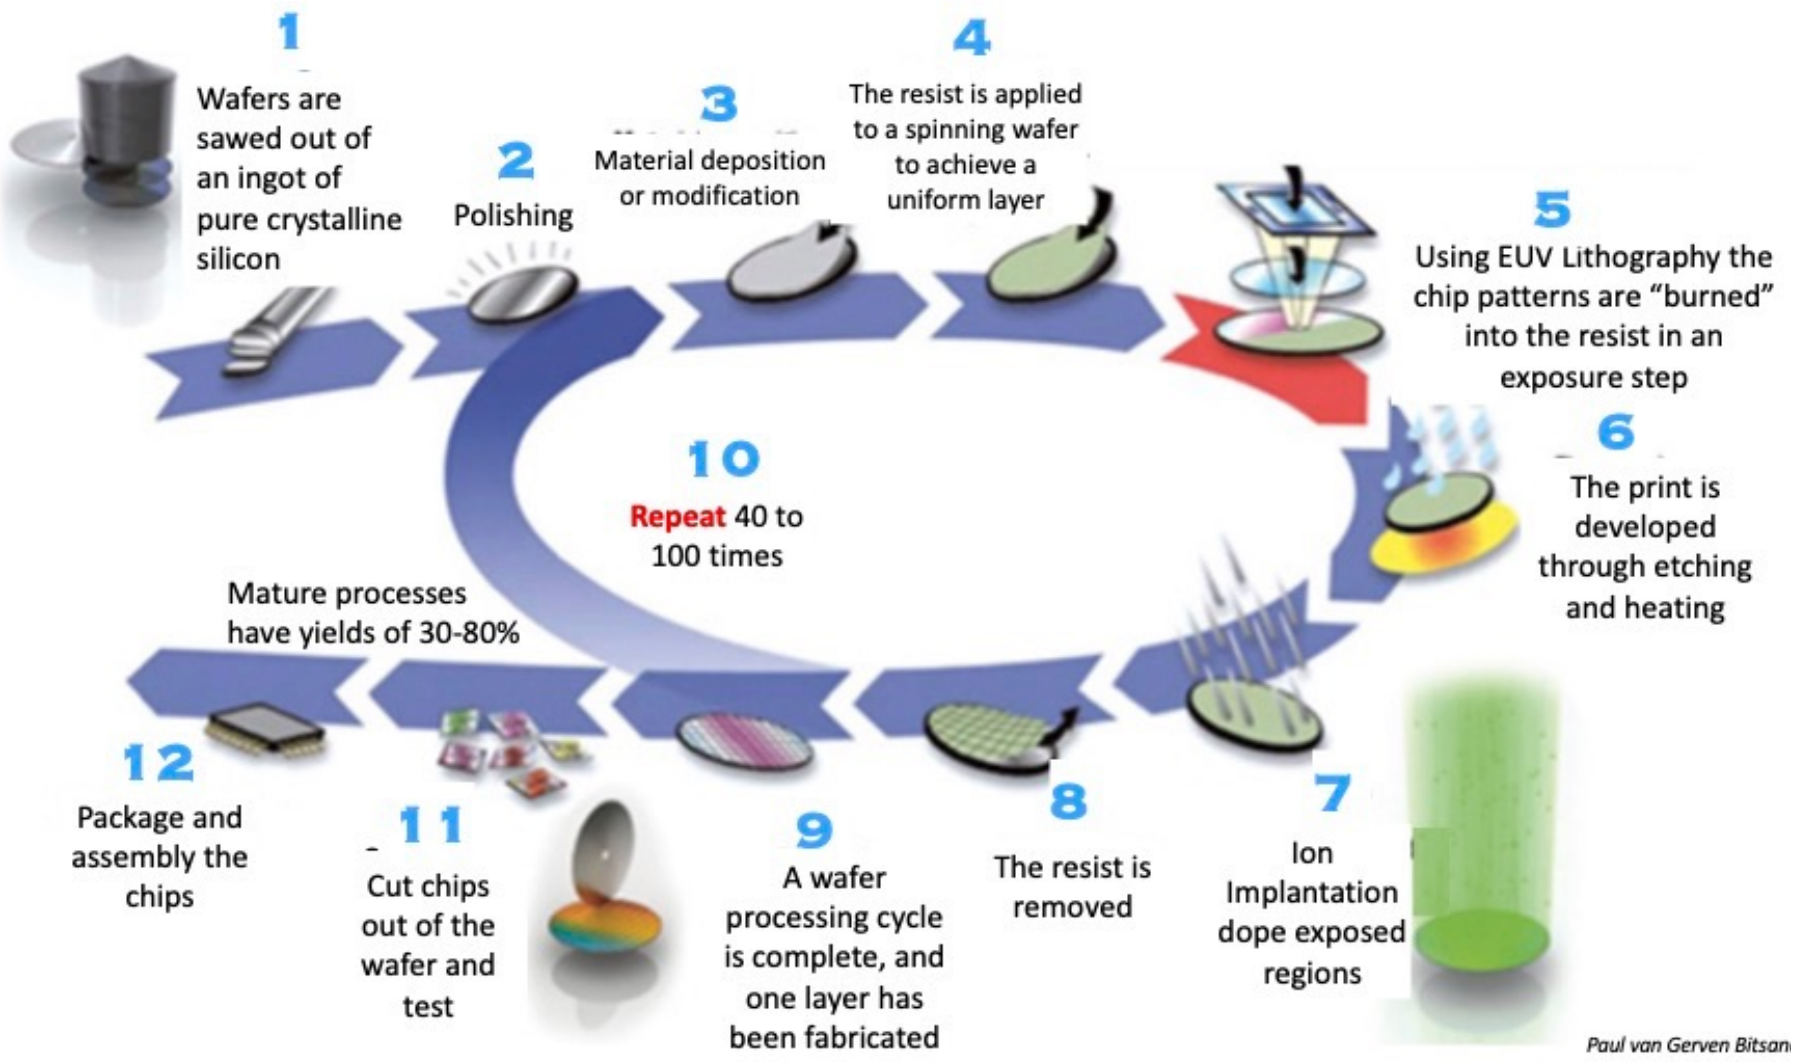

(Courtesy Alessandro Piovaccari)

Source: Bits & Chips

The Same Effort to Make 1 vs. 10B+ Transistors

- Expose entire chip in one shot
- Pattern smaller features by using shorter wavelengths of light to expose photoresist (historically started at 436nm, now 13.5nm)
- Commonplace today to print nm-scale patterns
- Must be done in “clean room” free of dust particles that will destroy integrity of mask patterns

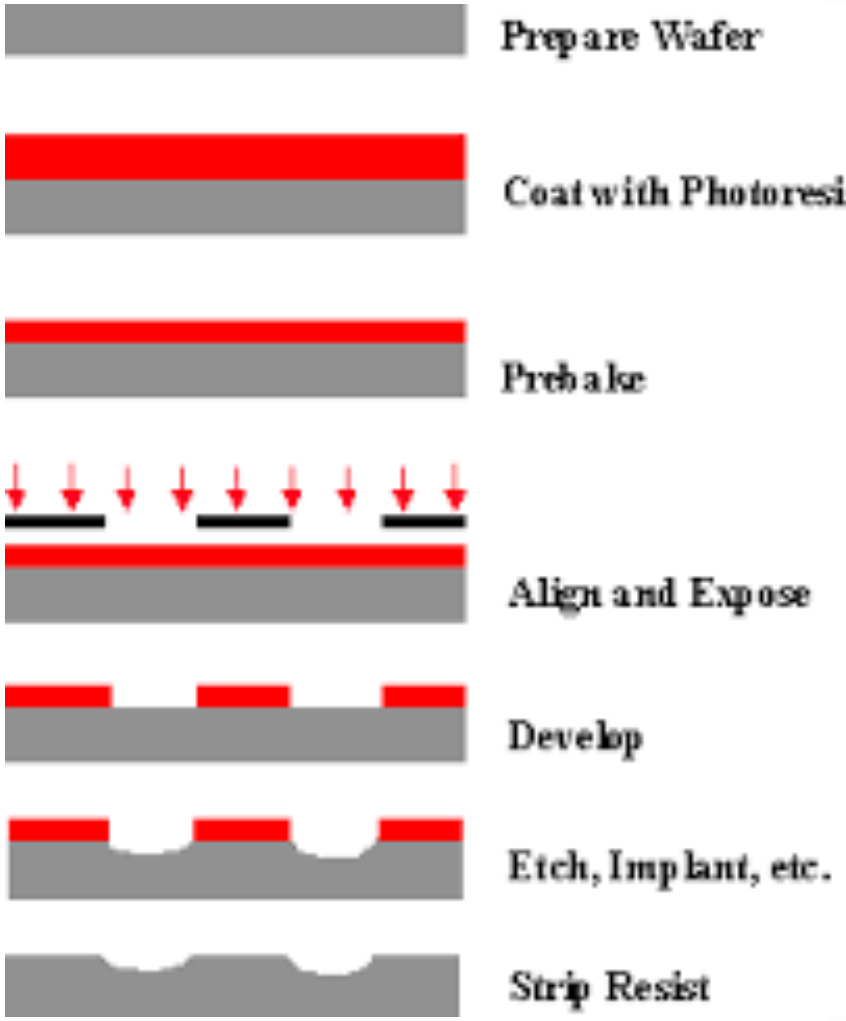

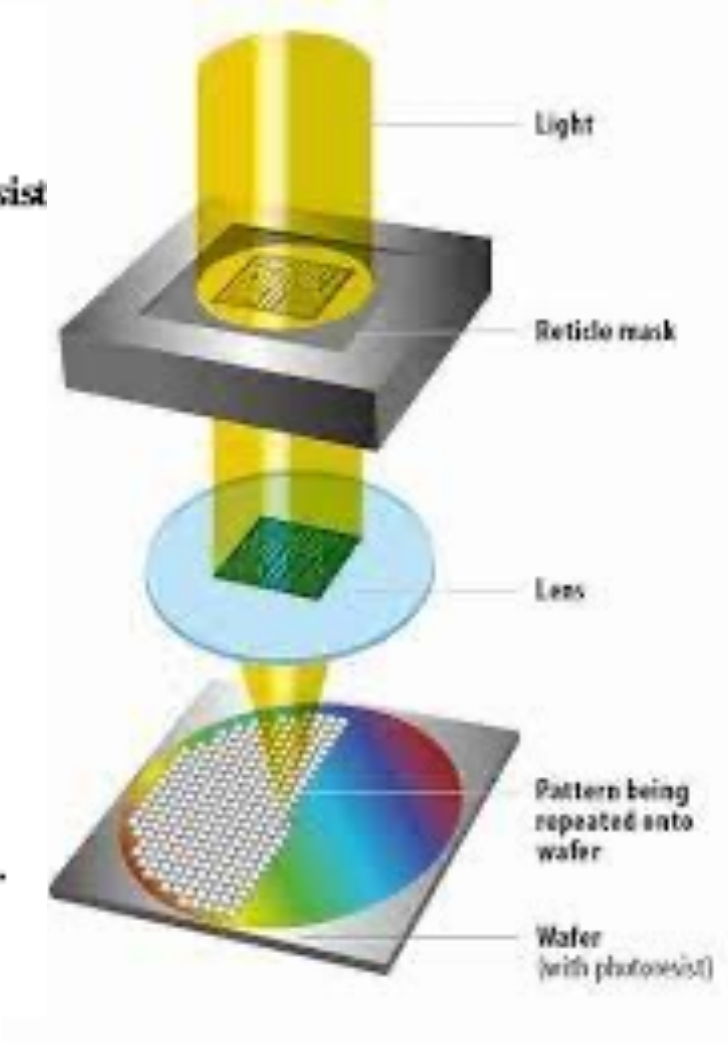

The Result

Top View of a Die

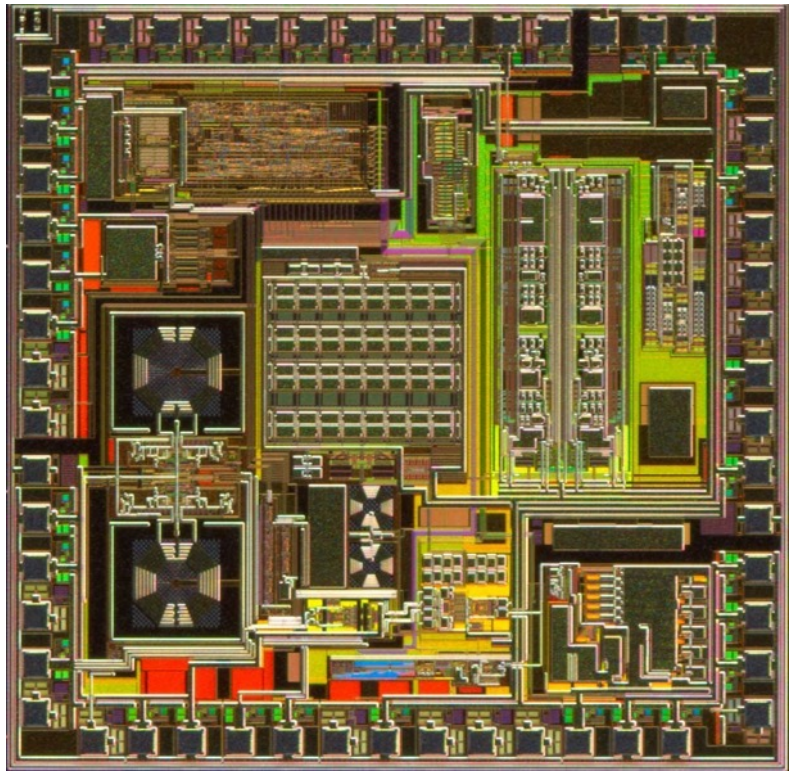

Source: IBM

Cross-section

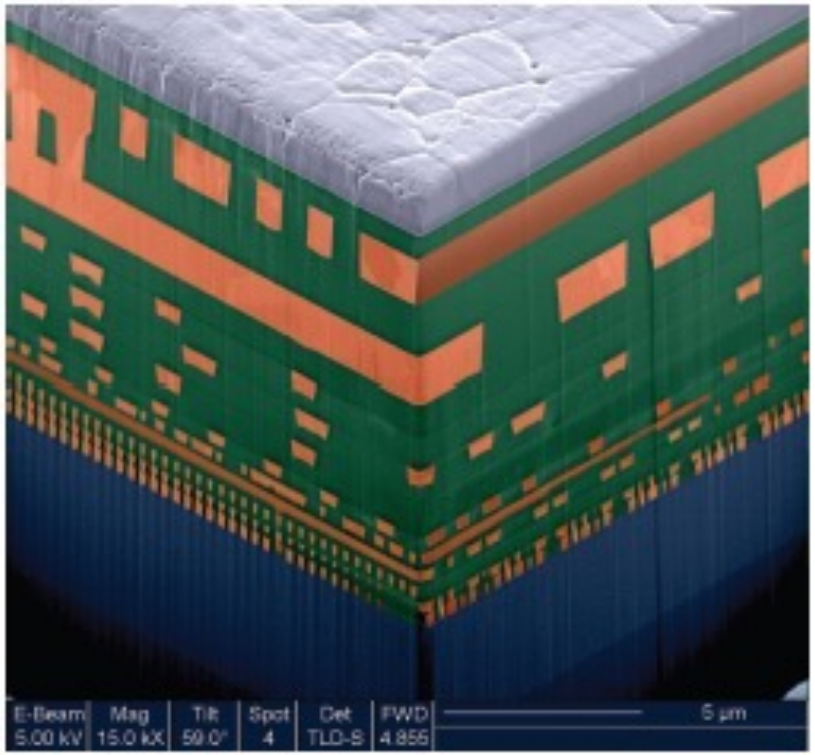

Source: IBM

Metal layers

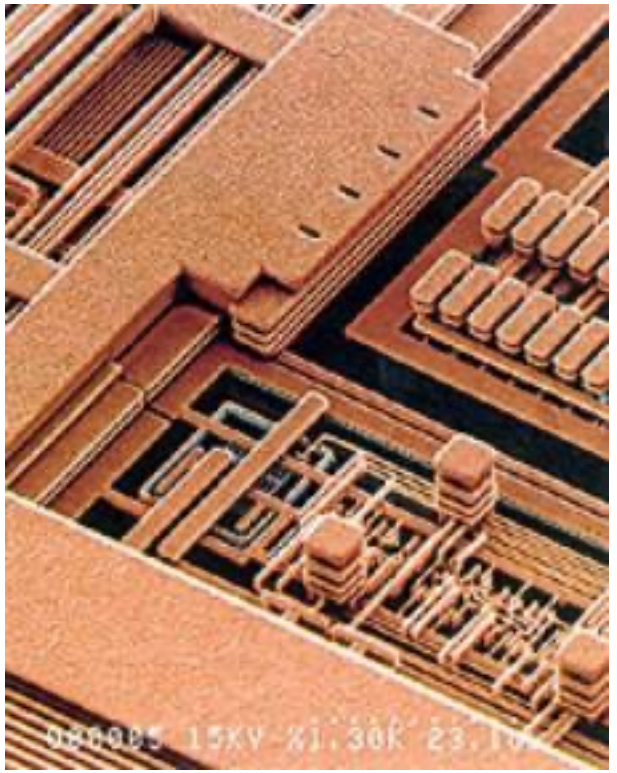

Source: IBM

(Courtesy Alessandro Piovaccari)

What's on a Chip?

- Billions of independent little capacitors, each of whose voltage represents stored information (e.g., logic 0 or 1)
- Capacitors charge/discharge (pull-up to supply or pull-down to ground) depending on logical function of inputs
- Switches implemented using transistors

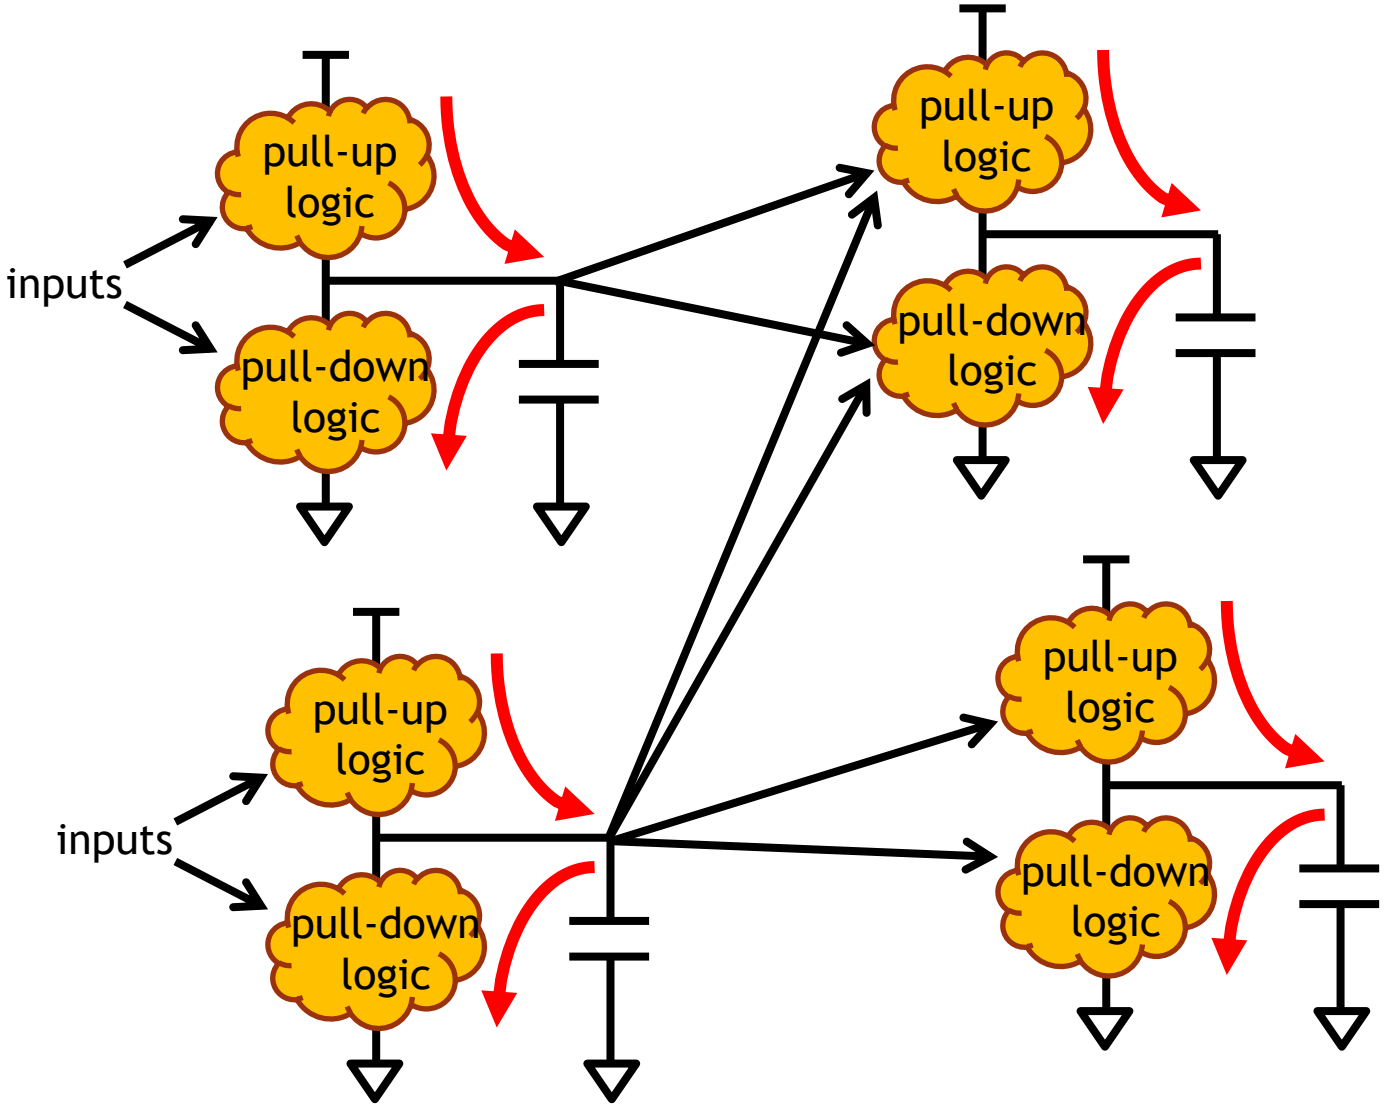

The Basics of Semiconductors

Pure semiconductor is an electrical insulator (not very interesting)

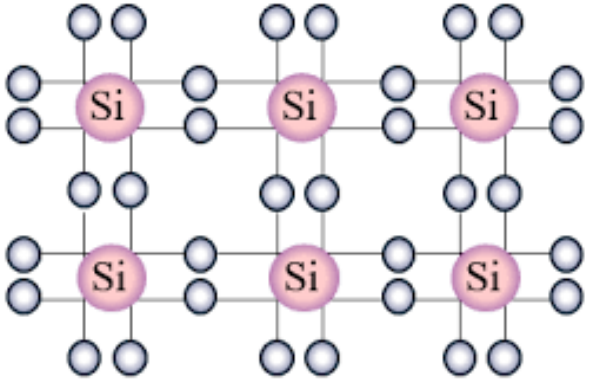

Replace occasional silicon atoms with impurity atoms (dopants) makes semiconductor an electrical conductor

n-type semiconductor

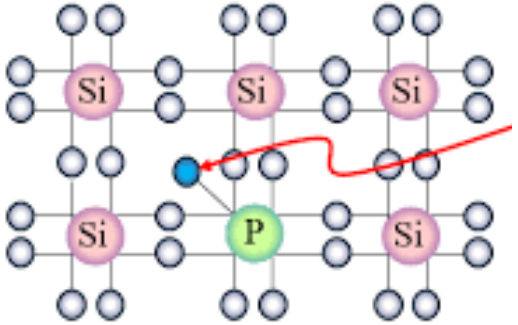

electrons can move around

p-type semiconductor

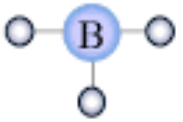

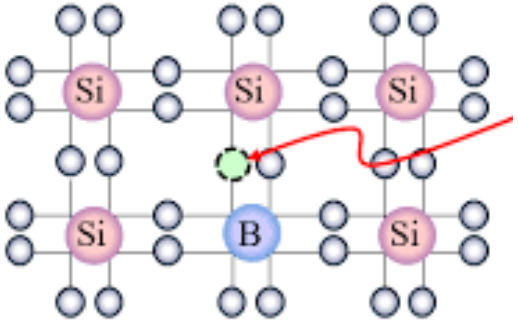

holes can move around

Source: shindengen.com

pn Junction (aka Diode)

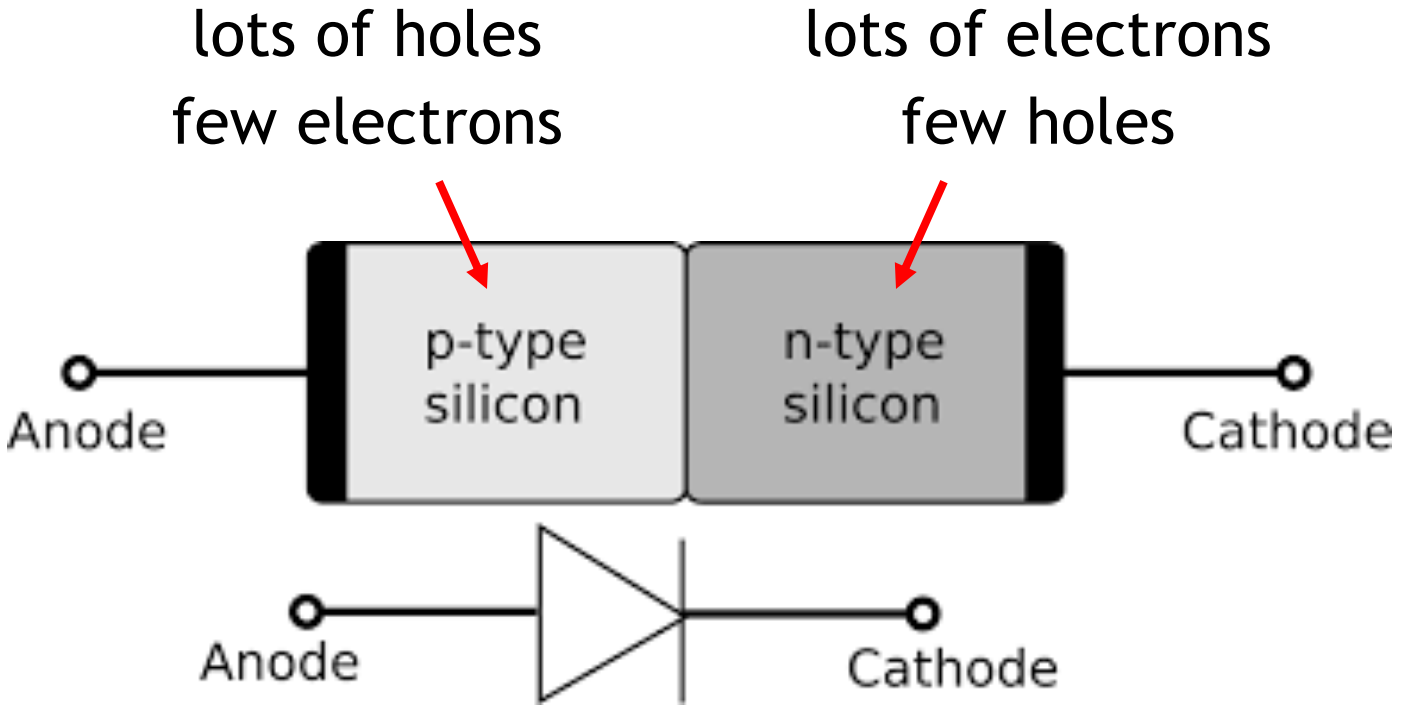

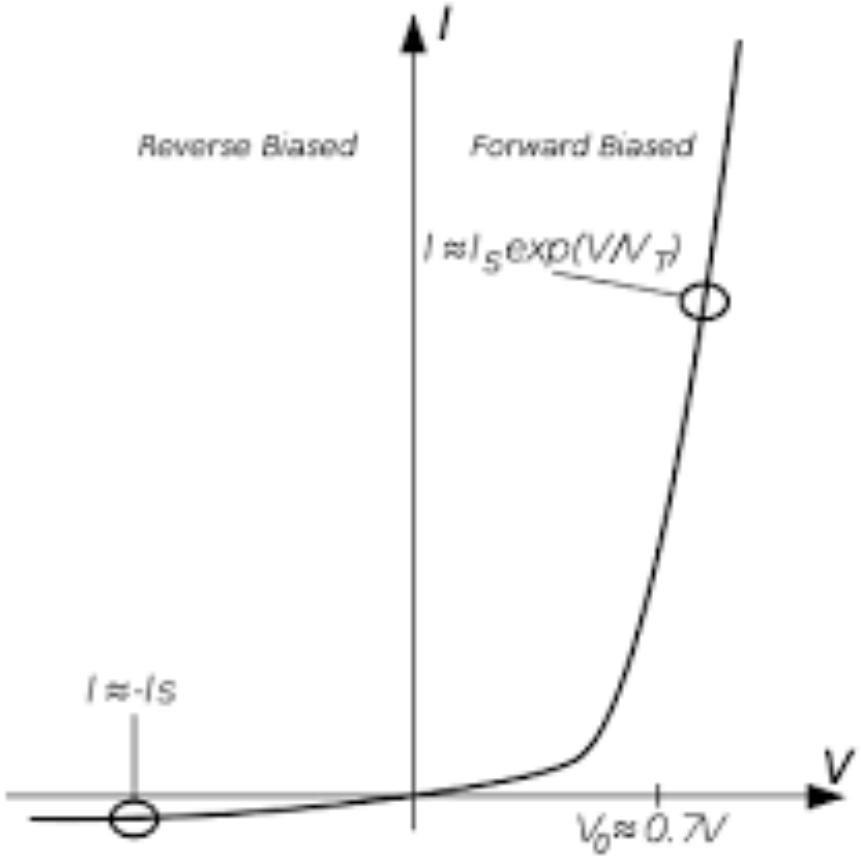

Allows current to flow
in only one direction

The MOSFET

- Metal-Oxide-Semiconductor Field Effect Transistor
- Patented in 1929, demonstrated in 1961
- Gate voltage controls current from source to drain
- Two flavors: N-MOSFET & P-MOSFET → Complementary MOS or CMOS
- Gate length is generally name of technology node (e.g., 32nm)

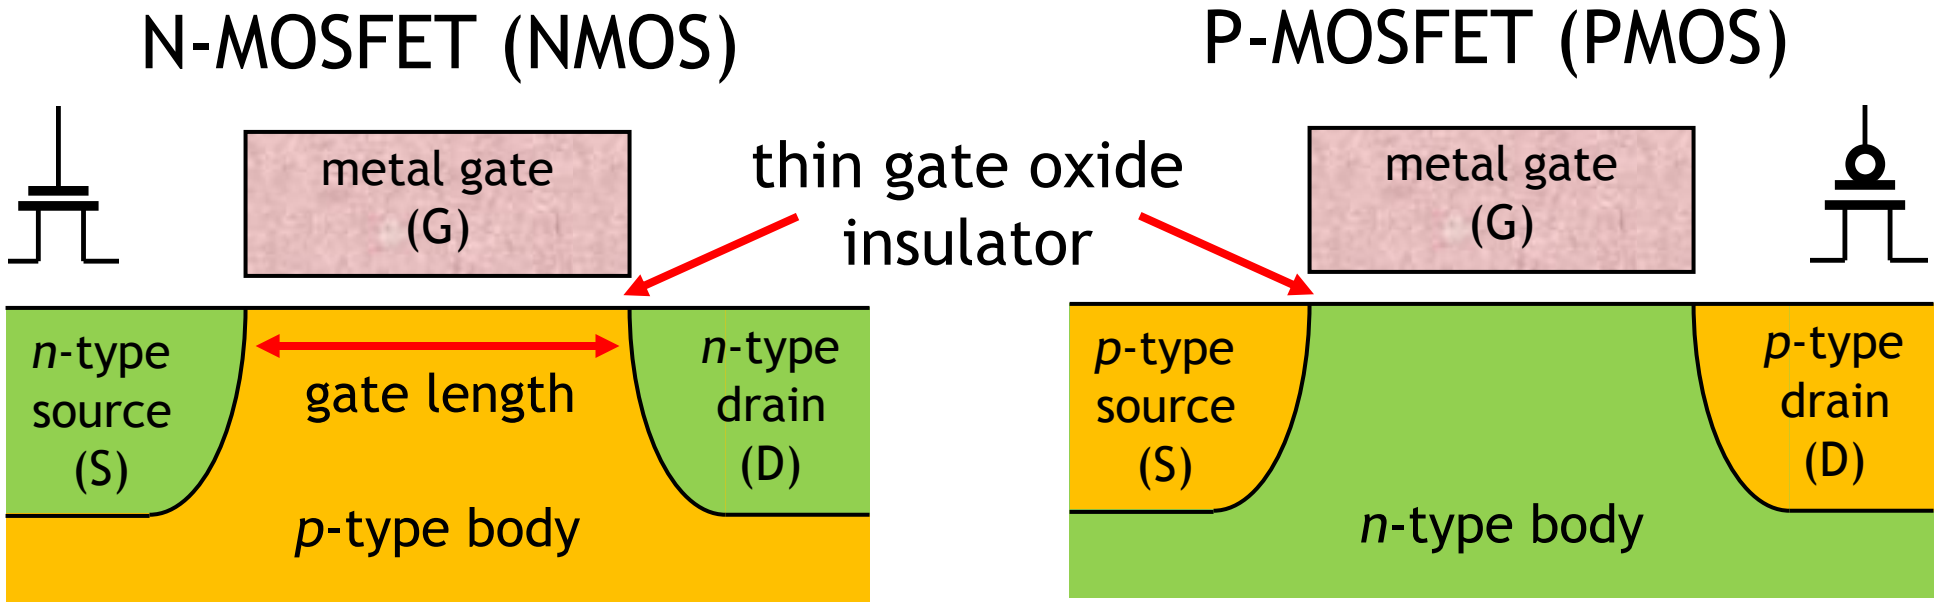

The N-MOSFET (NMOS)

No gate voltage.

What happens when you apply voltage across source/drain?

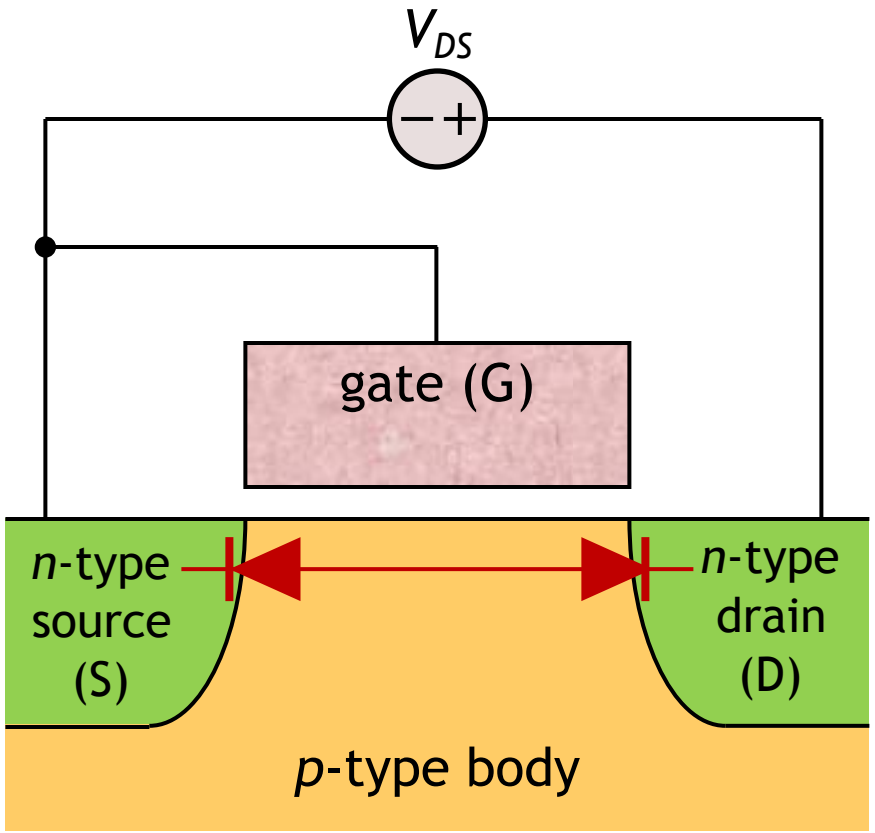

Apply gate voltage to create *n*-type inversion layer of electrons

What happens now when you apply voltage across source/drain?

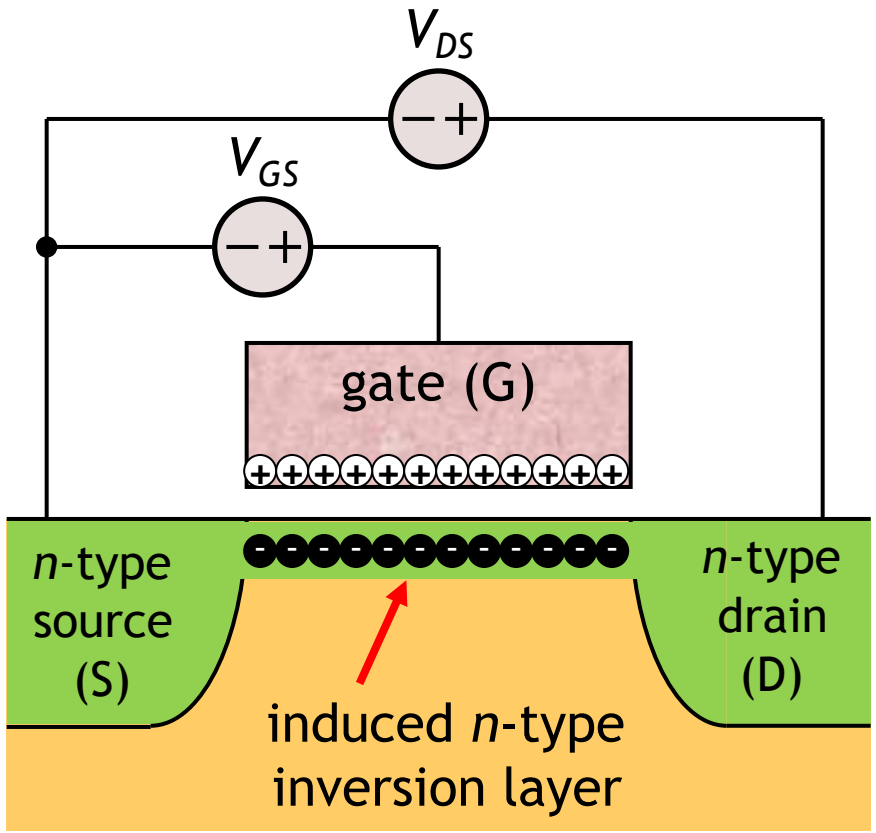

The P-MOSFET (PMOS)

No gate voltage.

What happens when you apply voltage across source/drain?

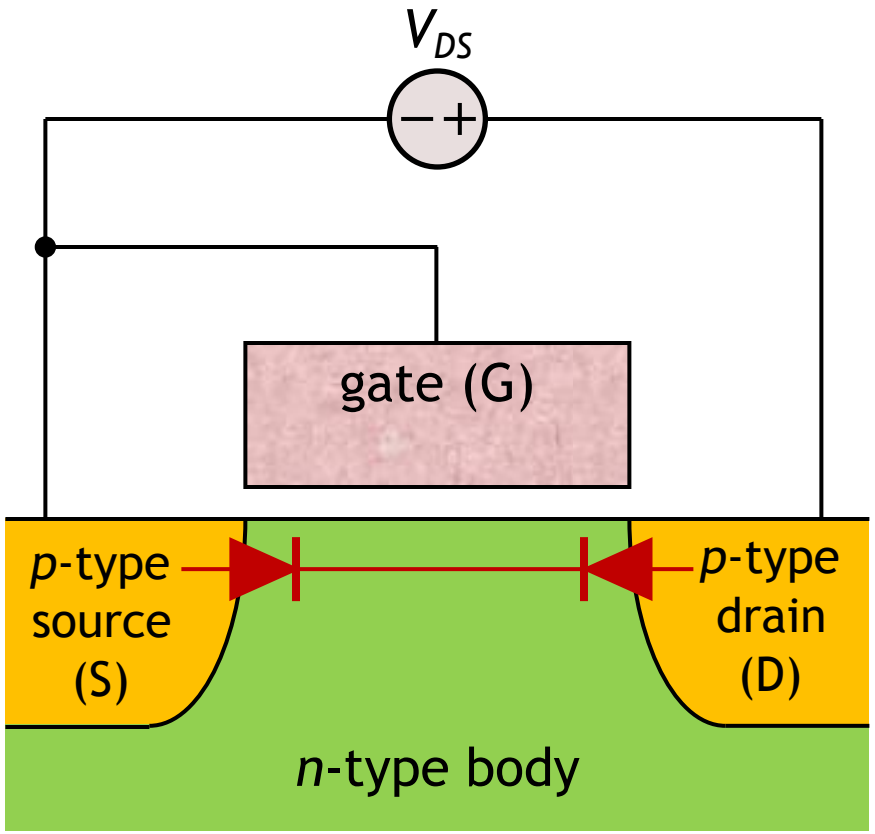

Apply gate voltage to create *p*-type inversion layer of holes

What happens now when you apply voltage across source/drain?

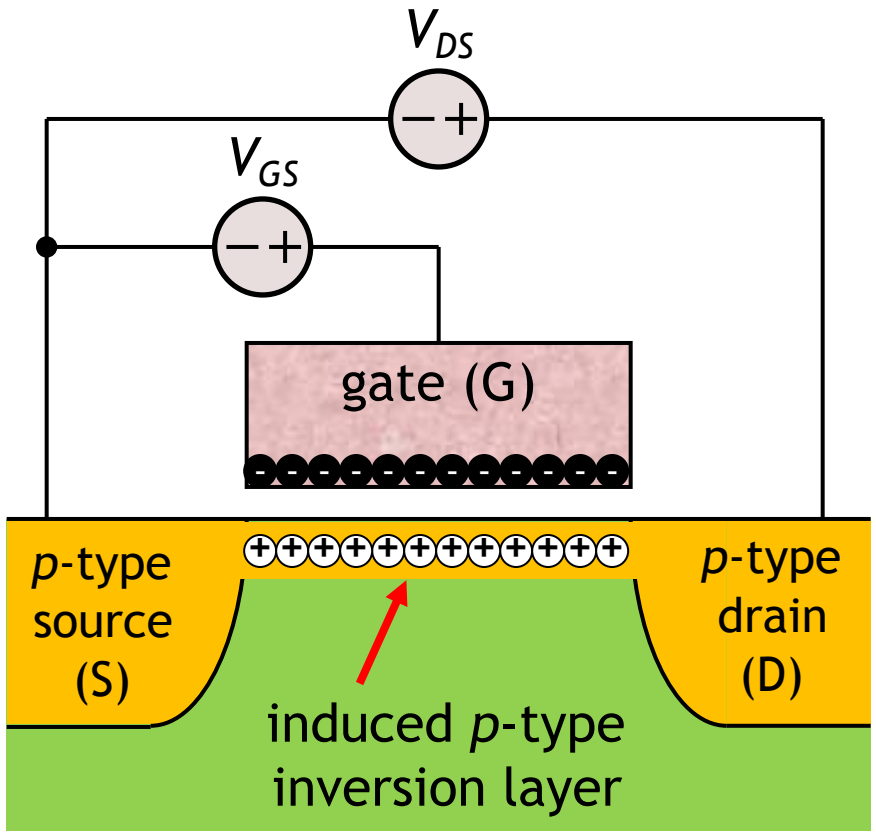

MOSFET I-V Characteristics

V_{TN} = NMOS threshold voltage (positive value)

V_{TP} = PMOS threshold voltage (negative value)

PMOS does not conduct current until $V_{GS} < V_{TP}$

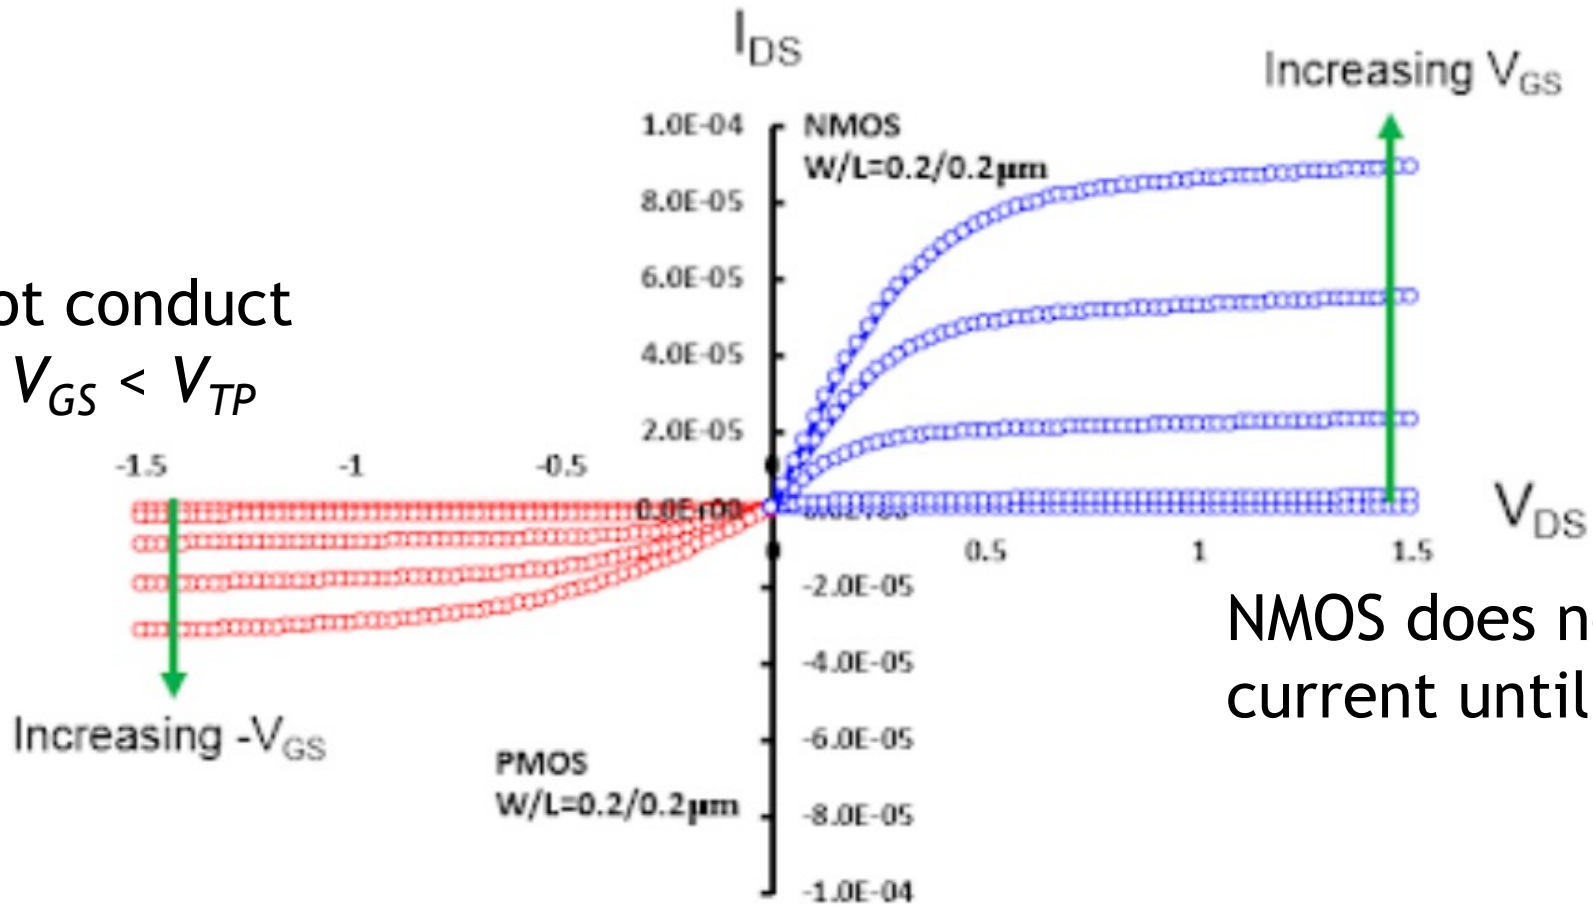

NMOS does not conduct current until $V_{GS} > V_{TN}$

The CMOS Inverter – Simplest Logic Gate

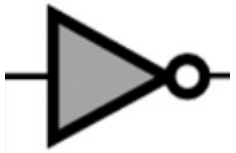

V_{in}	V_{out}
0	1
1	0

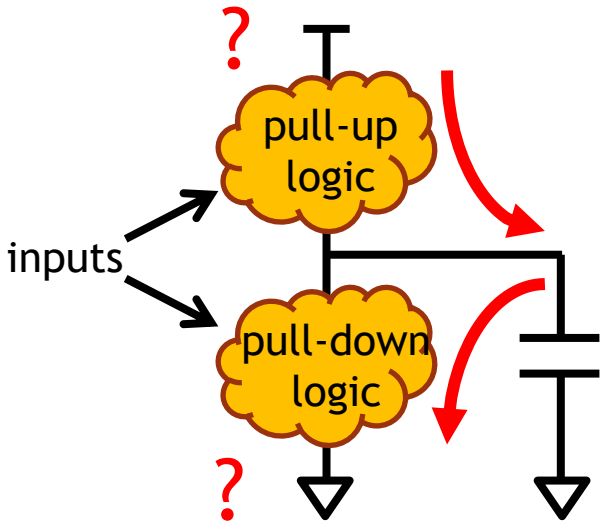

The Solution

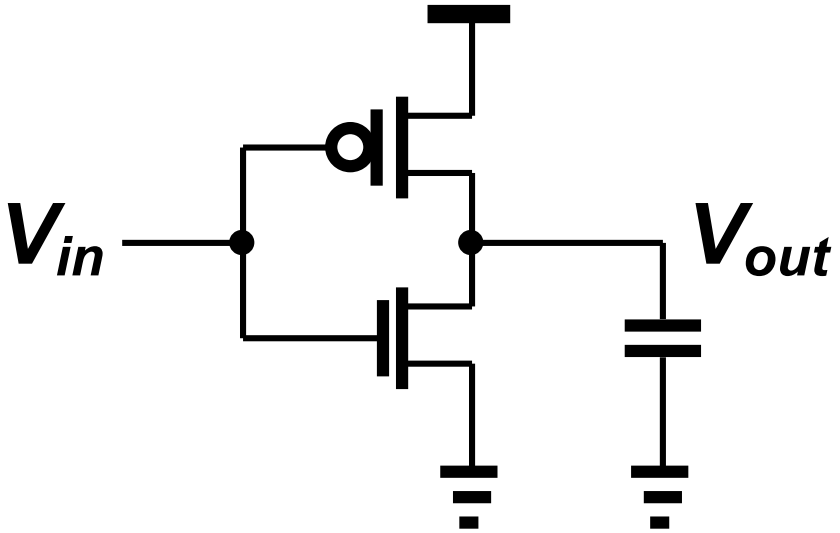

- V_{in}=0 → NMOS is off, PMOS is on → V_{out}=1
- V_{in}=1 → NMOS is on, PMOS is off → V_{out}=0
- No static current when input doesn't change → low power

Consider a Career in Semiconductors

The Opportunities

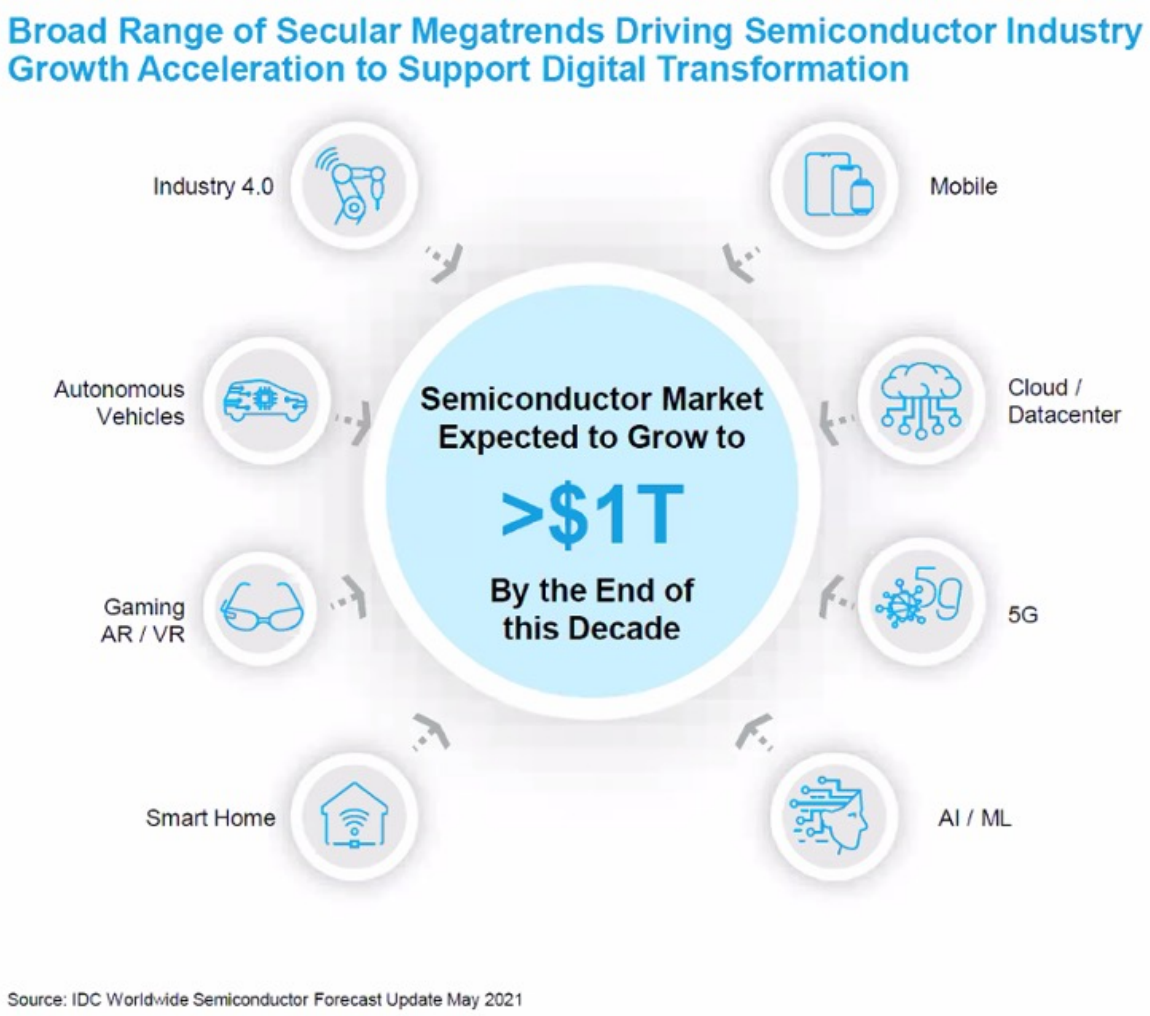

- Extremely multidisciplinary
  - Engineering (electrical, mechanical, materials, chemical, biomedical, systems, …)
  - Sciences (physics, chemistry, even biology)
  - Mathematics
  - Computer science (algorithms, AI)
- Very fast-paced, never a dull moment
- Change is normal
- Global
- Dress code is casual,
  otherwise I won't be employable ☺
- Pays well too ☺