CERN

Computing for Physics, yesterday today and tomorrow

Markus Schulz IT
markus.schulz@cern.ch

This is a very wide field...

- Computing is central for Physics and very diverse
- I’ll focus on the type of data processing that is done at the CERN computer centre
  - Skipping storage and networking

The Past

- The era of the mainframe (1958- 96)
  - Small number of really big machines
- The rise of the super computers (1980s →)
  - Very fast machines, many different concepts
- The workstations arrive (1990s)
  - Affordable machines linked together
- PCs can do it all (2000 →)
  - Many, many, many machines linked together

Mainframes at CERN

- Too many to mention all:
  - Ferranti Mercury
  - CDC 7600
  - IBM 370/168

Ferranti Mercury [1958-65]

- Vacuum tube machine
  - **60 microsec** clock cycle
    - 180000 times slower than a PC
    - Less complex operations (no division etc.)
  - Magnetic core memory (5 kByte)
  - Magnetic drum storage (80 kByte)
  - Paper tape I/0

Technology of the time

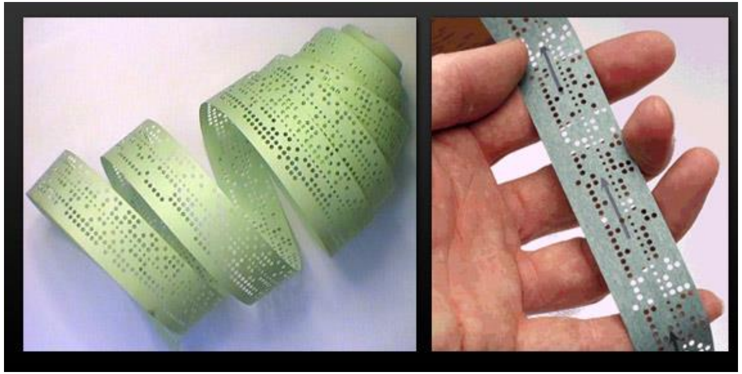

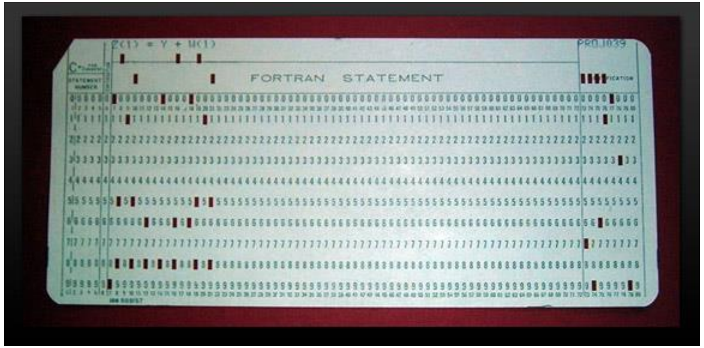

≈ 5 cm

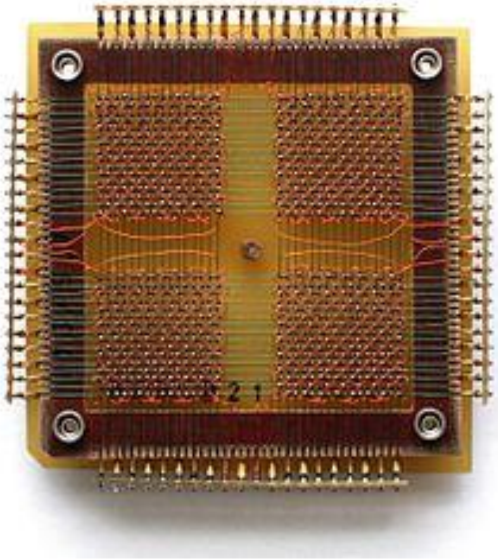

128 Byte

≈ 20cm

Tube based logic

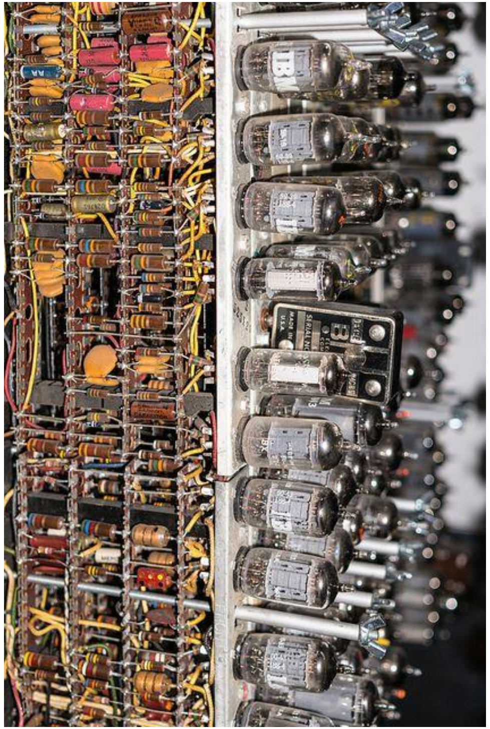

CDC 7600 [1972-84] ← 12.5 years

- Logic based on transistors, core memory
- 36 MHz clock rate (82 times slower than a PC)
- 480kByte fast memory, 3.7 MByte slow memory
- ≈ 40 Million Instructions Per Second (MIPS)
- Magnetic Tapes
- System of machines
- Reliability was an issue:
  - 1973: 3 hours between crashes
  - 1977: 16 hours

Technology

Magnetic Tape

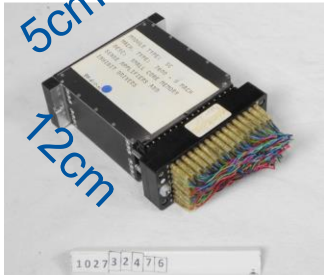

Module (core)

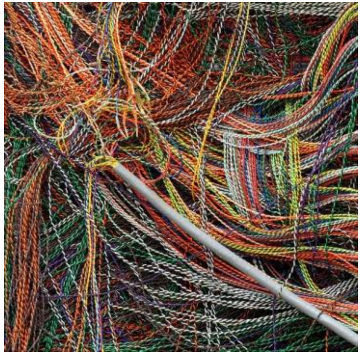

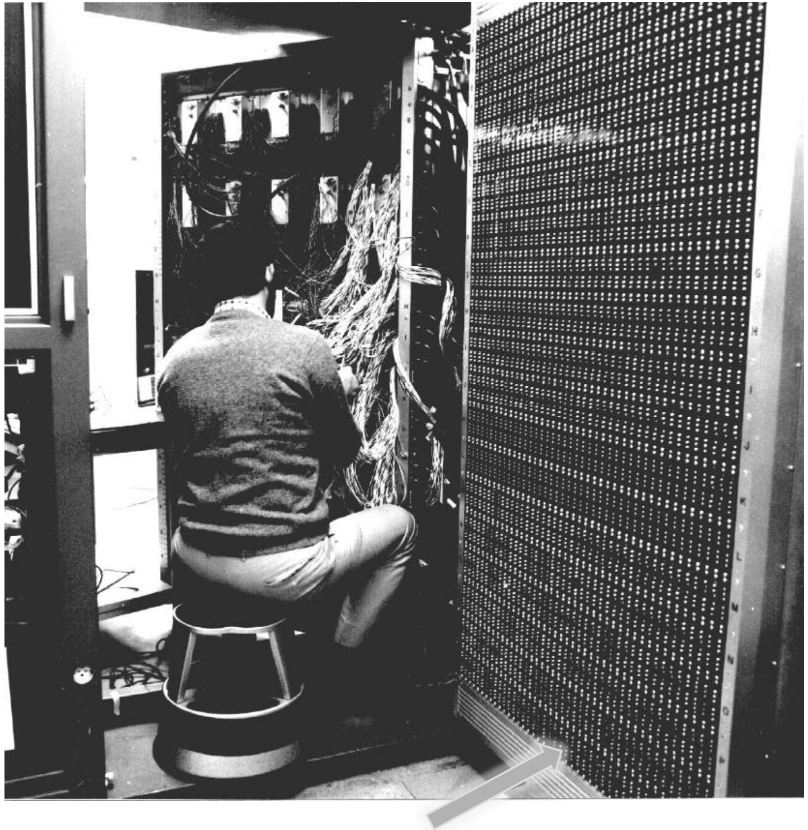

Modules

1978
Building 513

IBM 370/168 [1976-82]

- First machine with integrated chips and a modern architecture
  - Virtual memory, caches, robot tape system
  - Could support 200 users concurrently
- Semiconductor memory (4MBytes)
- 12 MHz (240 times slower than a modern PC)
- 3 MIPS
- The first of a long series of compatible IBM machines (3081, Siemens 7880, 3090/X)
  - CERN and SLAC build their own emulators

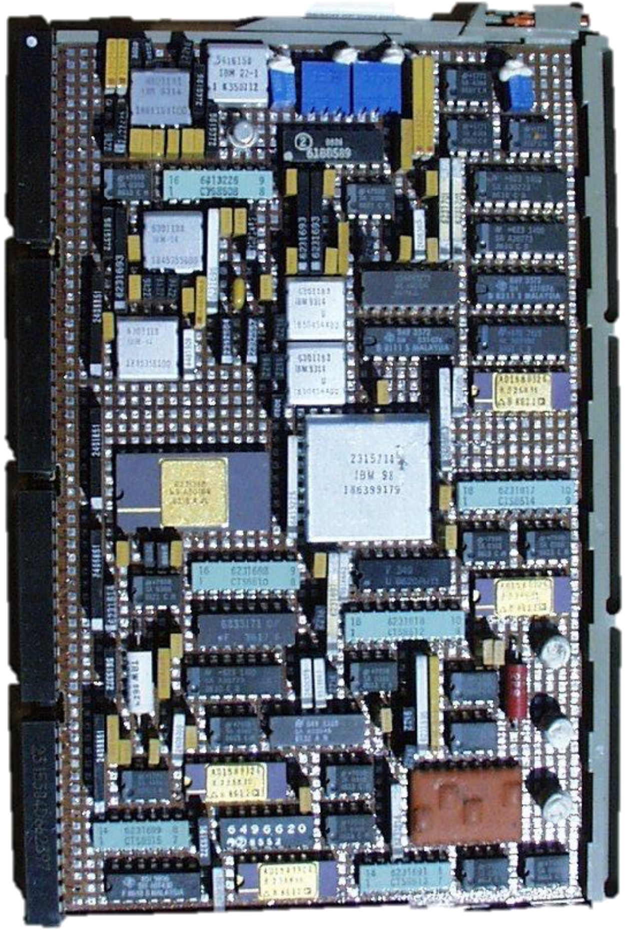

One of many cards

The Super Computers

- CRAY X-MP/48 [1988 – 1993]
- MEIKO CS-2 [1993-96]
  - 32 CPU system, high speed network

CRAY X-MP/48 [1988 – 1993]

- 4 CPUs 105 MHz (29 times slower than a modern PC)
  - Modern architecture with vector capabilities (SIMD)
  - 112 MIPS scalar power
  - Vector performance was much higher, but only few programs could profit from it
- 64 Mbyte fast memory and 1 GByte of slow memory
- 48 GBytes of disk
- Connected to IBM and VAX systems in the computer centre


ISSUE DE SECOURS
EMERGENCY EXIT
SSD
CRAY

The Workstations arrive

- In the 90s fully integrated CPUs with sufficient power became available
  - For relatively little money
- Many different architectures and vendors
  - HP/Apollo
  - SUN
  - SGI
  - IBM
  - DEC ....
- By linking many of them a lot of computing could be provided for a fraction of the cost

SHIFT [1990 – 200x]

- Scalable Heterogeneous Integrated FaciliTy
- Heterogenous system, UNIX variants as OS
  - Apollo, HP, SUN and SGI
  - The first 150.000 chf machine could deliver ½ of a CRAY X-MP
  - Grew steadily until it was replaced by commodity PCs in the early 2000s

CERN


HP - SHIFT
DELPHI - CMS
ALICE - NA49 - ATLAS
S U N
SHIFT - CMS
SGI - SHIFT
HP - CSF
CORE
ALEPH

The Reign of the Mundane

- IBM PCs became the de facto standard for desktop machines
- The price/performance has been unmatched
- First generation of PC clusters relied on desktop machines
  - Several disadvantages, density, cooling, reliability
- Second generation are based on rack mounted systems
  - Same technology basis, better mechanical and thermal properties
- Currently CERN runs more than 14.000 servers with 230.000 CPU Cores


First Generation

CERN Computing Power

- 1990 1000 Million FLOPS
- 2019 460 Tera FLOPS
- We now have 460000 times the computing power we had 29 years ago installed in the computer centre

The Future?

- All advancements have been driven by:
  - Technology evolution follows scaling “laws”
    - Moore’s laws (transistor size and price)
      - CPUs and memory
    - Similar laws for disks and networks
  - Markets
    - Investments follow the market

Single core performance is slowing down

- 52%/year from 1986-2003
- 3.5%/year from 2015
- But you get more cores for the same price!

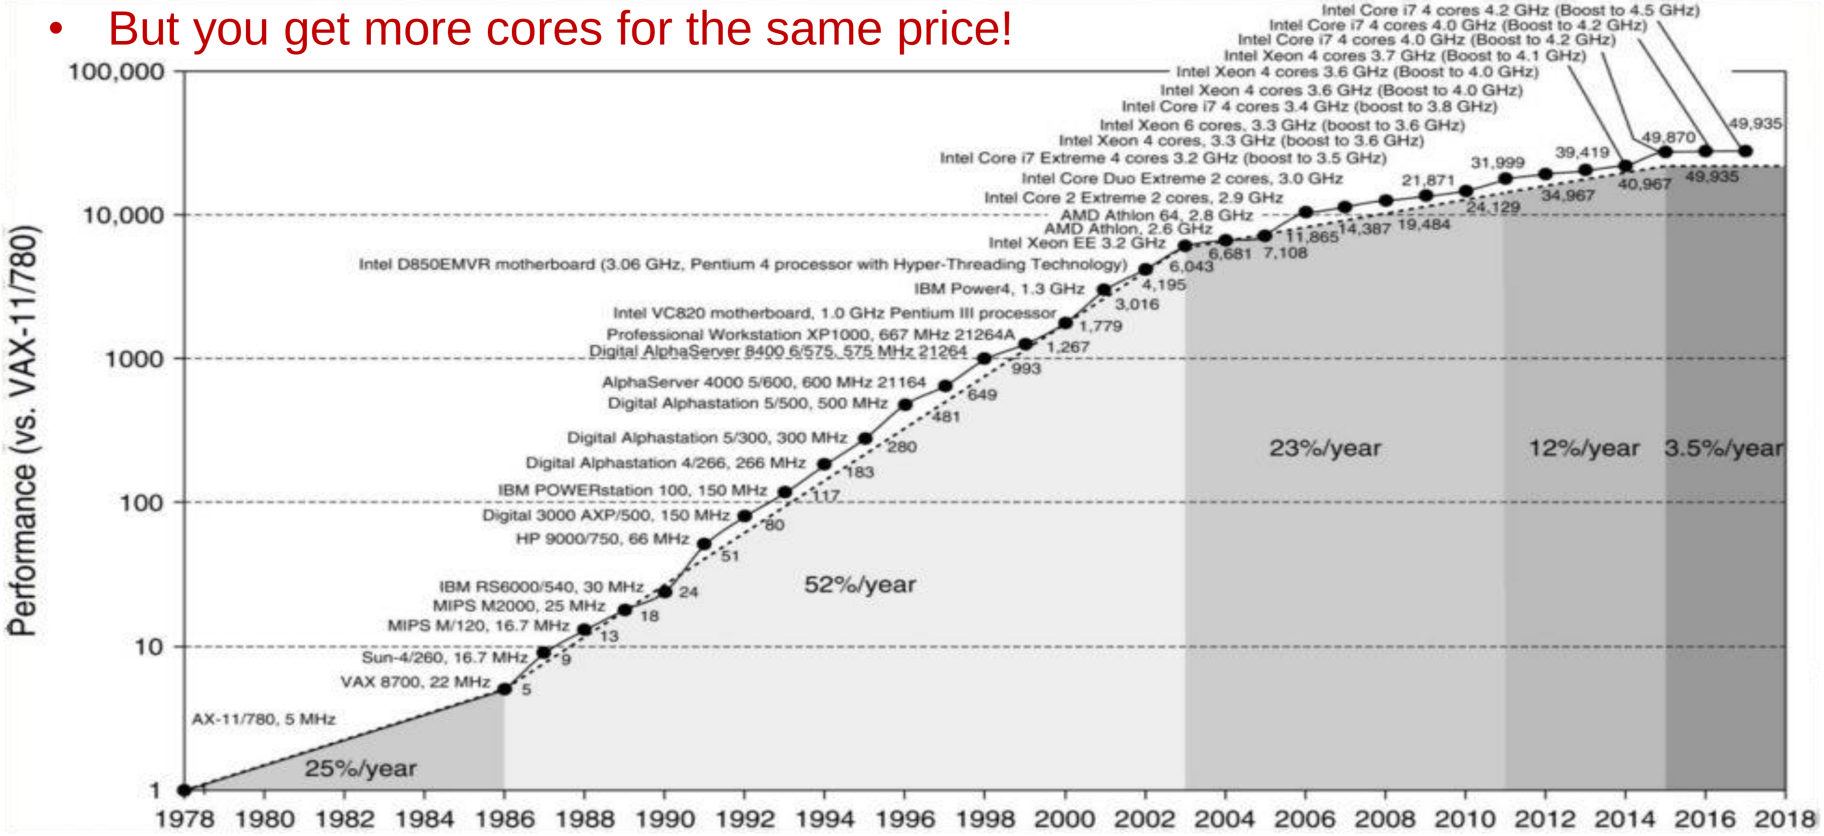

Will shrinking come to an end?

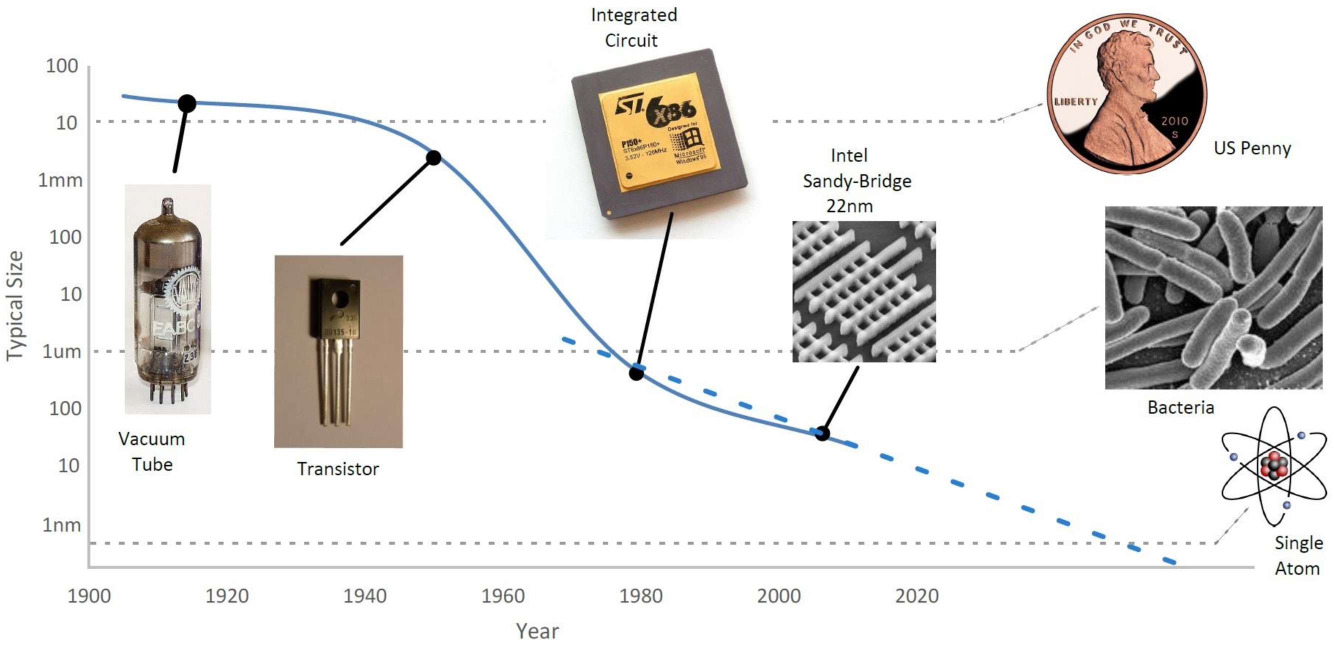

https://theglobalscientist.files.wordpress.com/2014/03/fig1.jpg

Yes! Long before reaching the size of an Atom

Four paths to more performance

- Parallel processing
  - Difficult to exploit
- Problem oriented hardware (accelerators)
- Alternative technologies
  - Many.......
- Quantum Computing
  - Still relatively young field, but remarkable progress and massive investments

What will it all mean for Physics?

- 5-10 years
  - 40% increase of single core performance
  - 5 times the number of cores
  - Huge challenges for the software (parallel programming)
- 20 Years
  - Systems will reach the molecular level
    - Quantum effects will dominate → difficult to predict
- 50 Years
  - Either something boring, or something marvellous will happen ;-)

Thanks a lot

www.cern.ch

And enjoy the (next) paradigm shift !!!! ☺

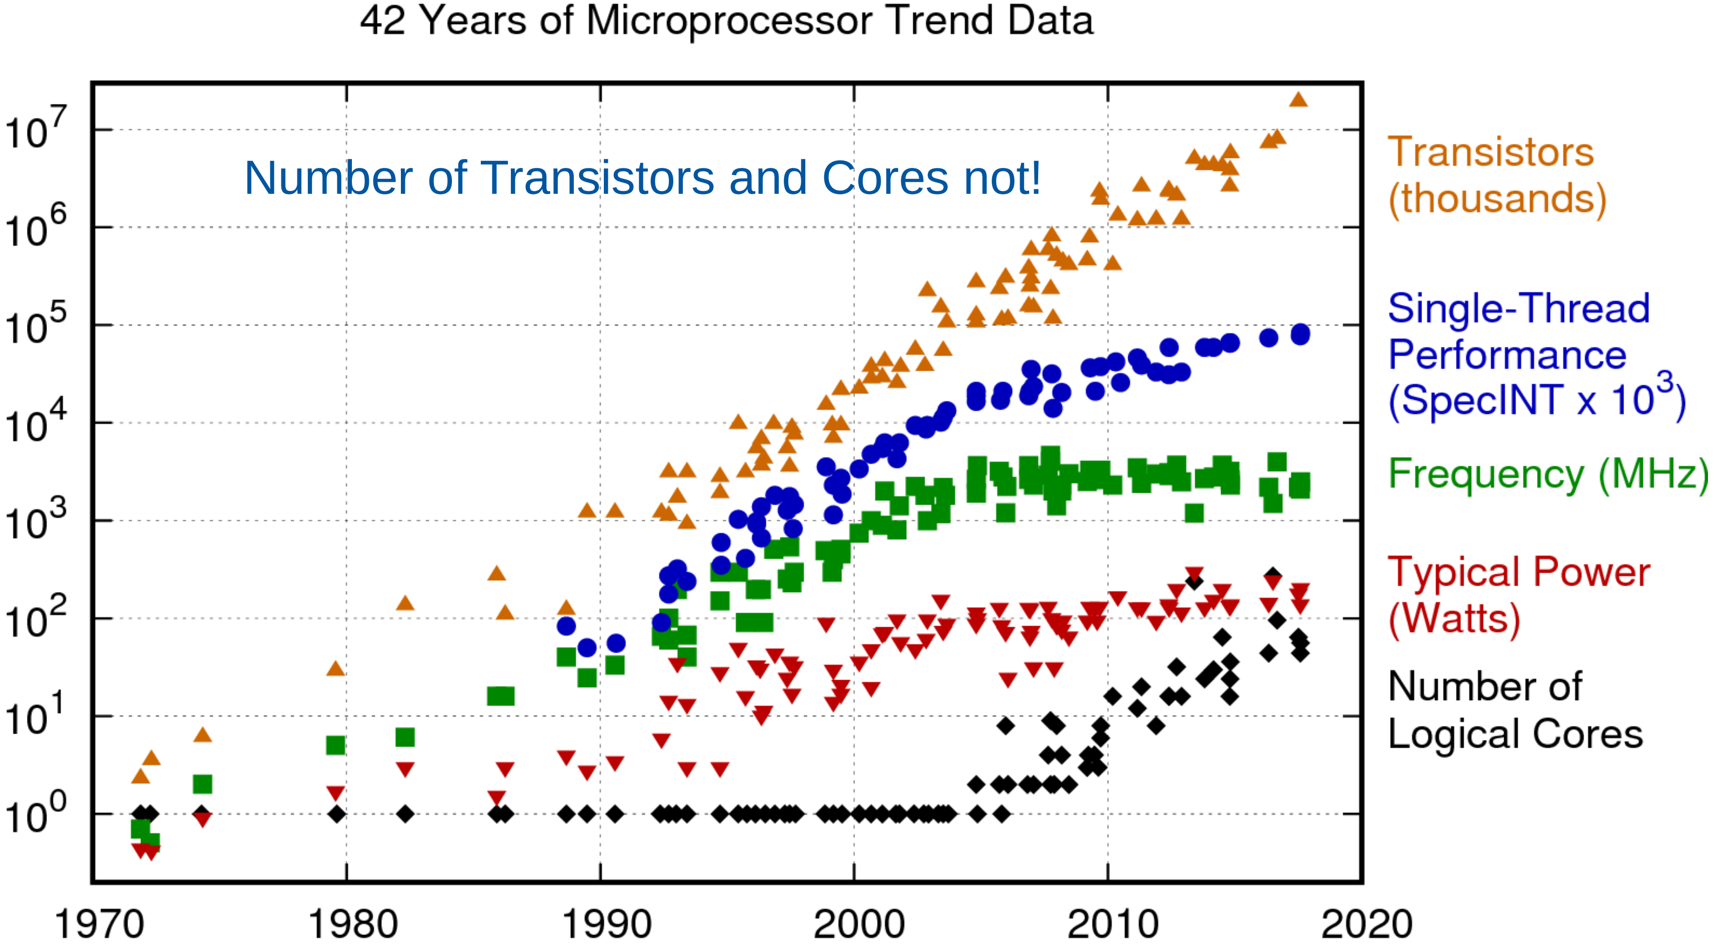

Original data up to the year 2010 collected and plotted by M. Horowitz, F. Labonte, O. Shacham, K. Olukotun, L. Hammond, and C. Batten
New plot and data collected for 2010-2017 by K. Rupp

Markets

- Consumers focus shifts to price and power
  - Smart phones, IOT, etc.
- Number of producers decreases
  - INTEL > 60% of processor market

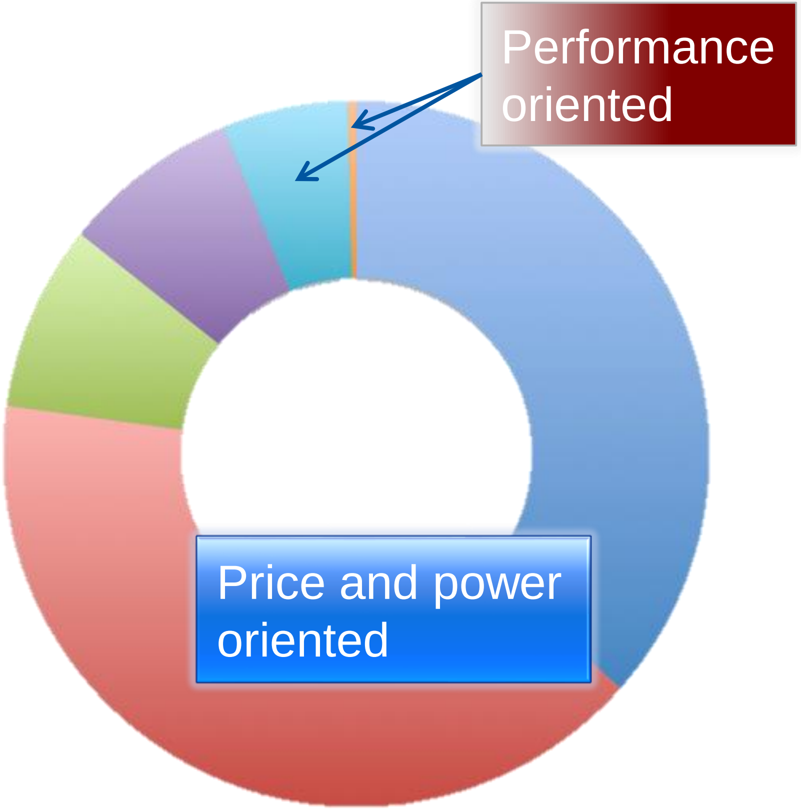

Alternative Technologies

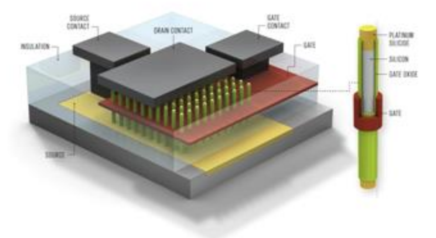

- Nanotubes, Molecular transistors, DNA, Graphen, .
  - proof of concepts (1-1000 components)
  - many challenges ahead, but promising
- 3D Structures
  - difficult to handle the heat
  - on the surface already widely used
- Memristors, high density, first products soon, but minuscule compared to discs
- Laser nanostructures in glass → 360 TB glass disk
- DNA for storage
  - proof of concept (1cm^3 == 1000 TB)
  - very, very expensive, slow