CMOS – Chip Level

Last updated 2/25/21

CMOS – Chip Level

- Millions to hundreds of millions of gates
  - Physical Space
  - Sheer numbers of gates
  - Keeping shared resources close to multiple users (memory)
  - I/O pin access and placement
  - Interconnect – getting all the wires connected
    - Typical processes have 6 – 10 layers of interconnect
    - Cell, local, global, power
  - Performance
  - Power / Heat Dissipation

CMOS – Chip Level

- Performance Drivers
  - Process Technology
    - Transistor performance
    - Short channel vs. long channel devices
    - High Vt and Low Vt devices
  - Clock Frequency
    - Maximum is set by the longest unit delay
    - Very complex timing tools used to ensure max frequency
  - Interconnect
    - RC delays
    - Capacitive coupling

CMOS – Chip Level

- Power Drivers
  - Process Technology - Dynamic, Static, Short circuit (D/S/SC)
  - Number of gates – D/S/SC
  - Clock Frequency – D/SC
    - Dynamic power becomes CV^2f, where f is clock frequency
    - Short circuit power is also multiplied by f
  - Supply Voltage – D/S/SC
  - Routing Efficiency – D/SC
    - Minimizing capacitance is critical

CMOS – Chip Level

- Power / Performance Balance
  - Device Level Solutions
    - Multiple T_{ox} (V_T)devices
      - Thin T_{ox} devices for high performance paths
      - Thick T_{ox} devices
        - Lower leakage
        - Higher voltages
  - Chip Level Solutions
    - Reduced interconnect R and C
    - Power islands – gating the power to circuits not in use
    - Clock Gating (module level)
      - Turn off the clocks to circuits not in use
      - Requires synchronization

CMOS – Chip Level

- Power / Performance Balance
  - System Level Solutions
    - Dynamic Voltage Scaling – changing VDD as needed
    - Dynamic Frequency Scaling – changing the clock frequency as needed
    - Together these are referred to as DVFS
  - Architectural Solutions
    - Pipelining
    - Multi-core processors
      - Homogeneous – dual/quad core
      - Heterogeneous – big/little/GPU
    - Memory Hierarchy

CMOS – Chip Level

- Power – Performance Tradeoffs

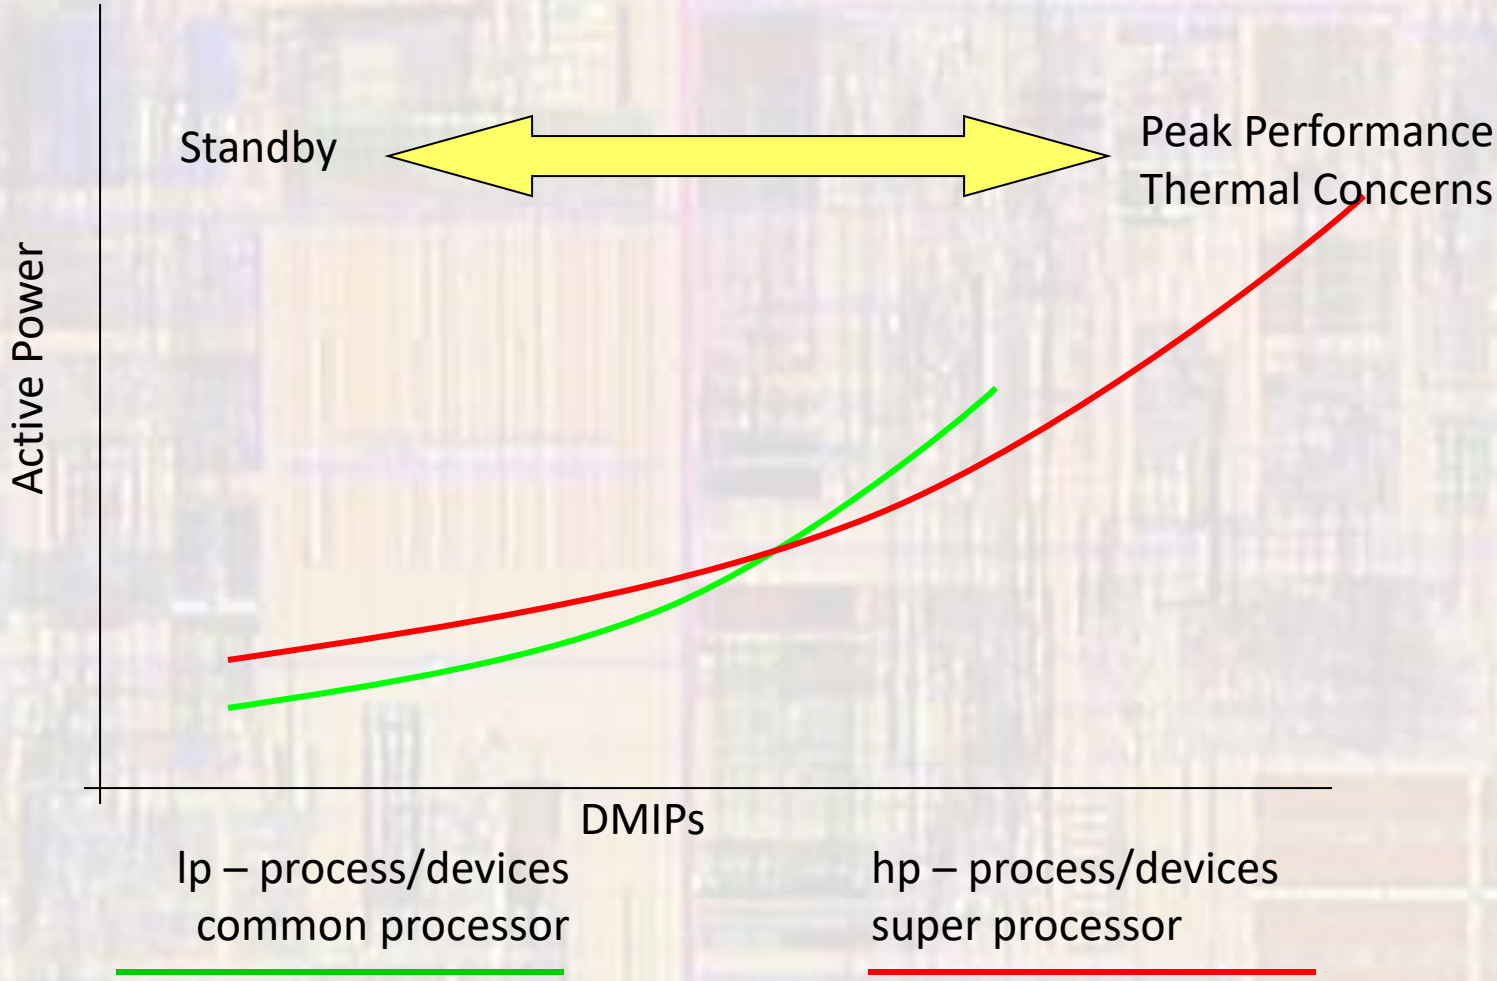

CMOS – Chip Level

- Power – Performance Tradeoffs

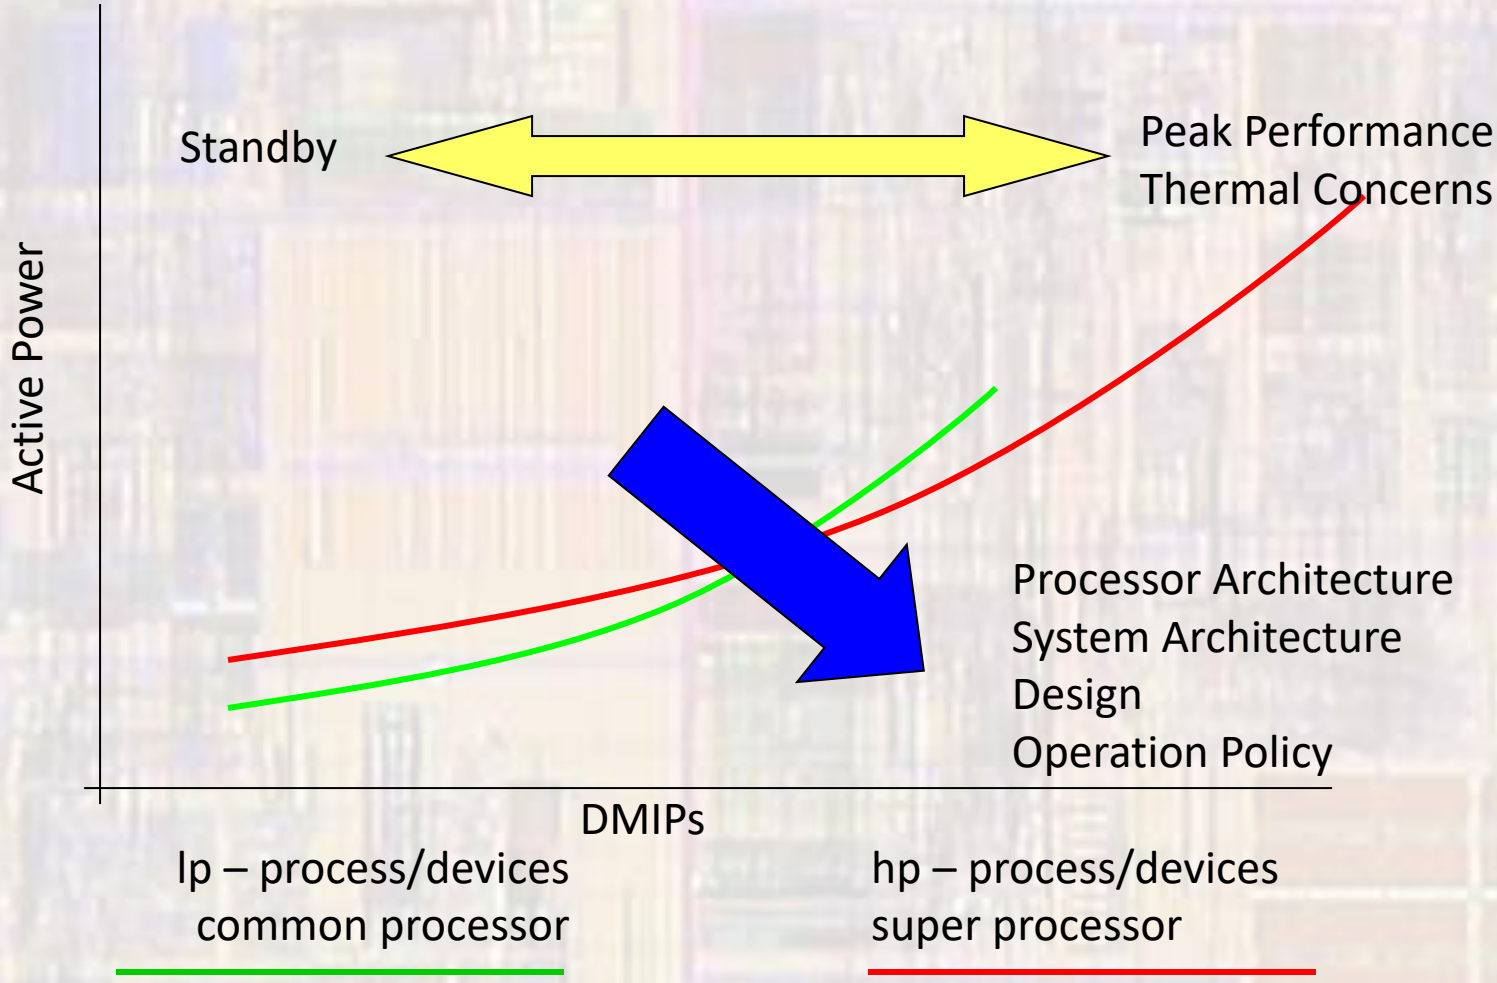

Technology

- Processor Performance

IEEE IRDS

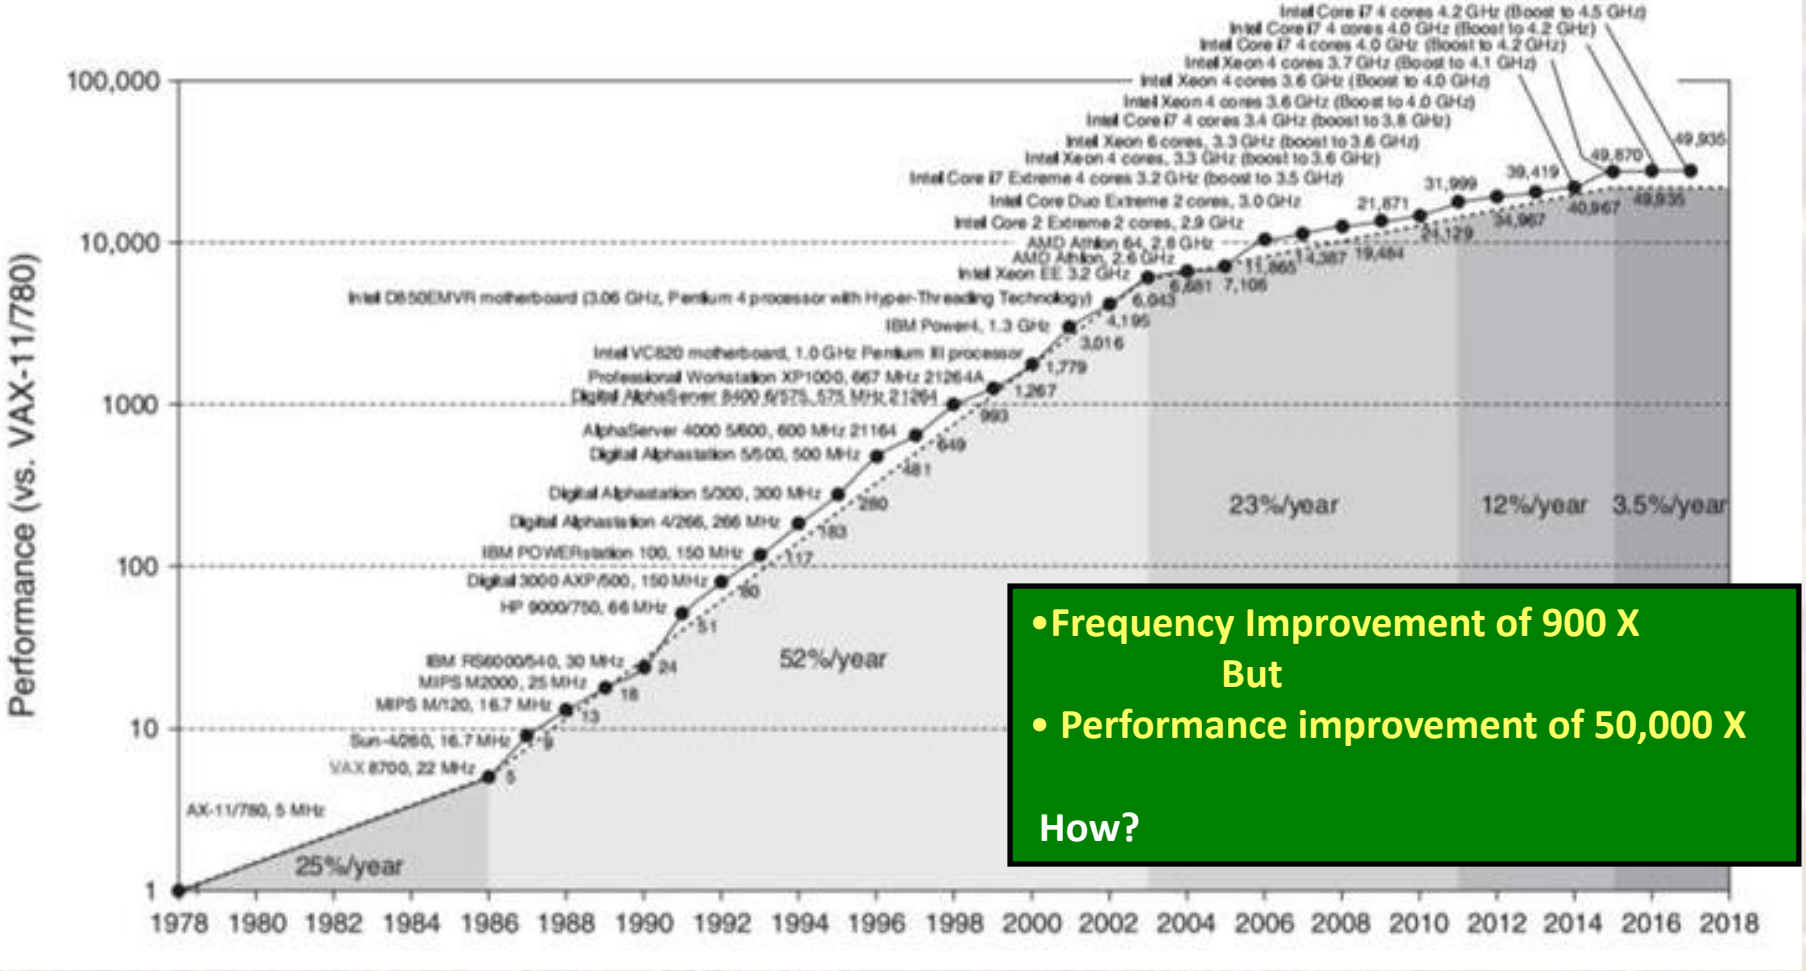

Source: Computer Architecture, Hennessy and Patterson, 2017 Elsivier Inc

CMOS – Chip Level

- Processor Performance

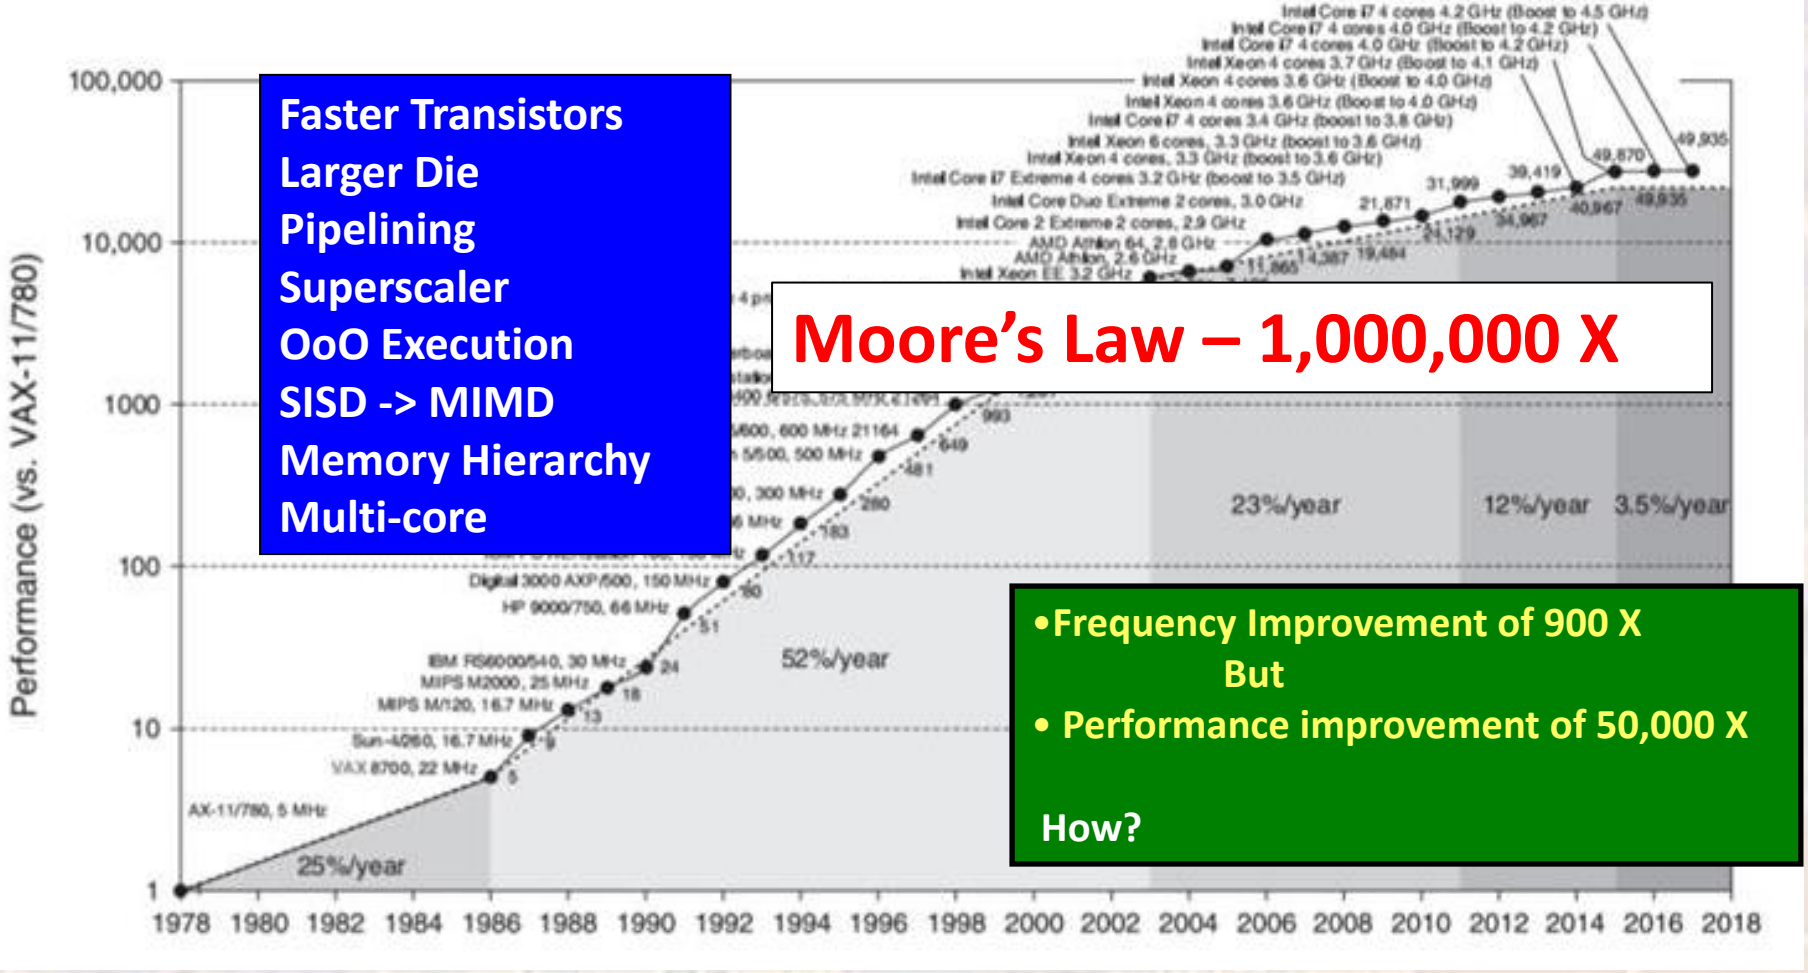

Source: Computer Architecture, Hennessy and Patterson, 2017 Elsivier Inc

CMOS – Chip Level

- Transistors per chip

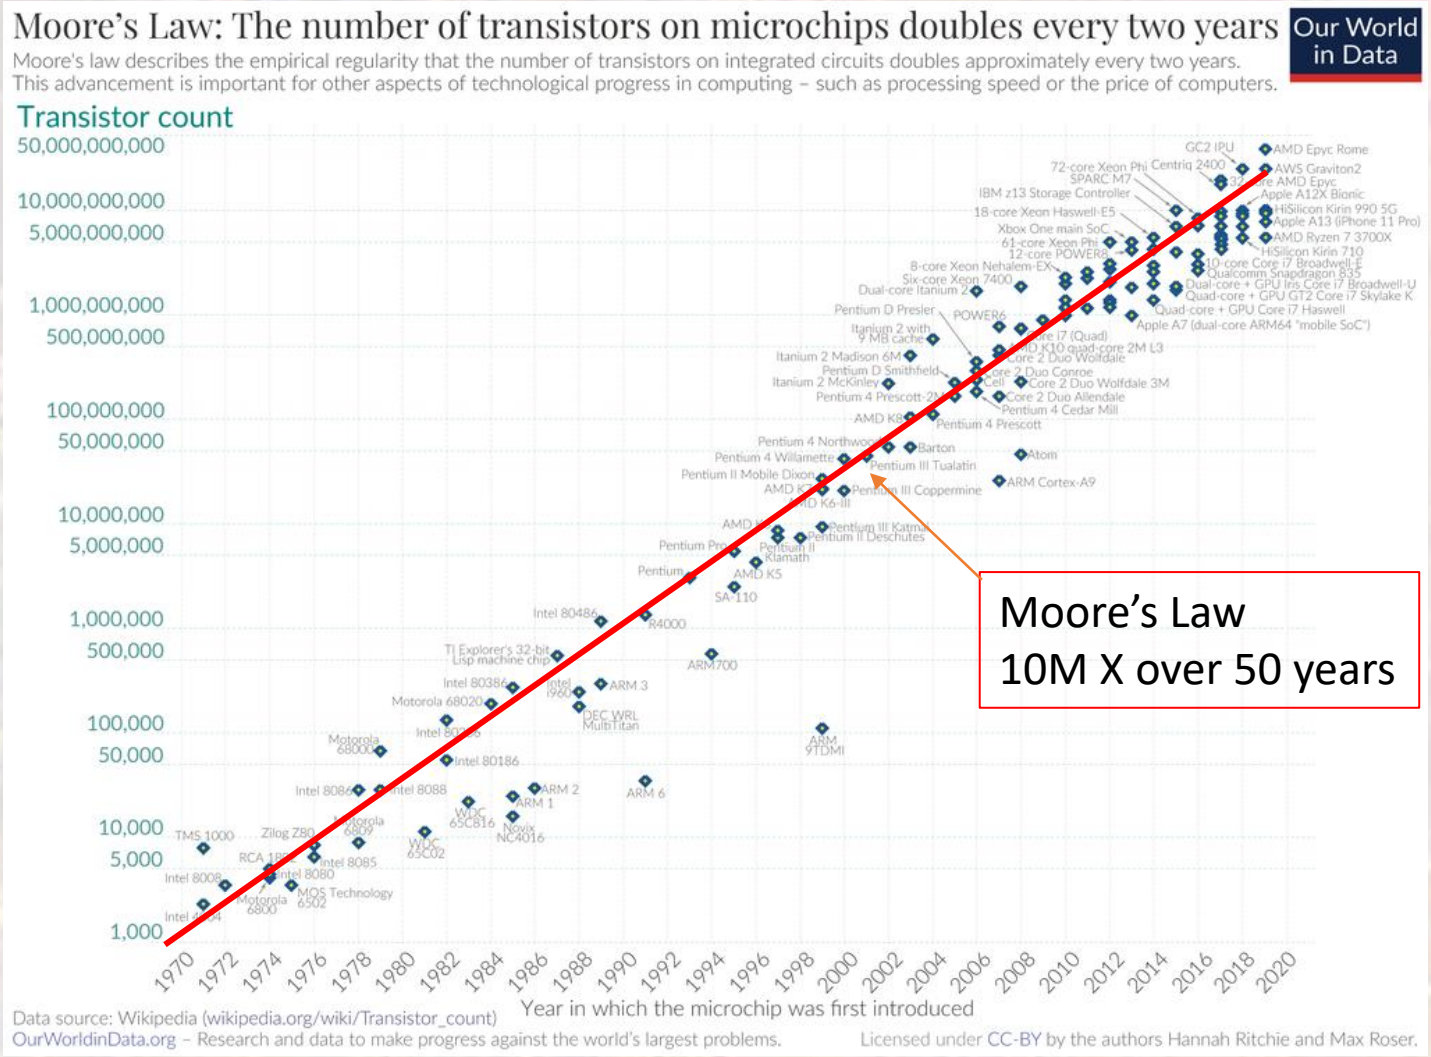

CMOS – Chip Level

- SPECint Performance – single CPU

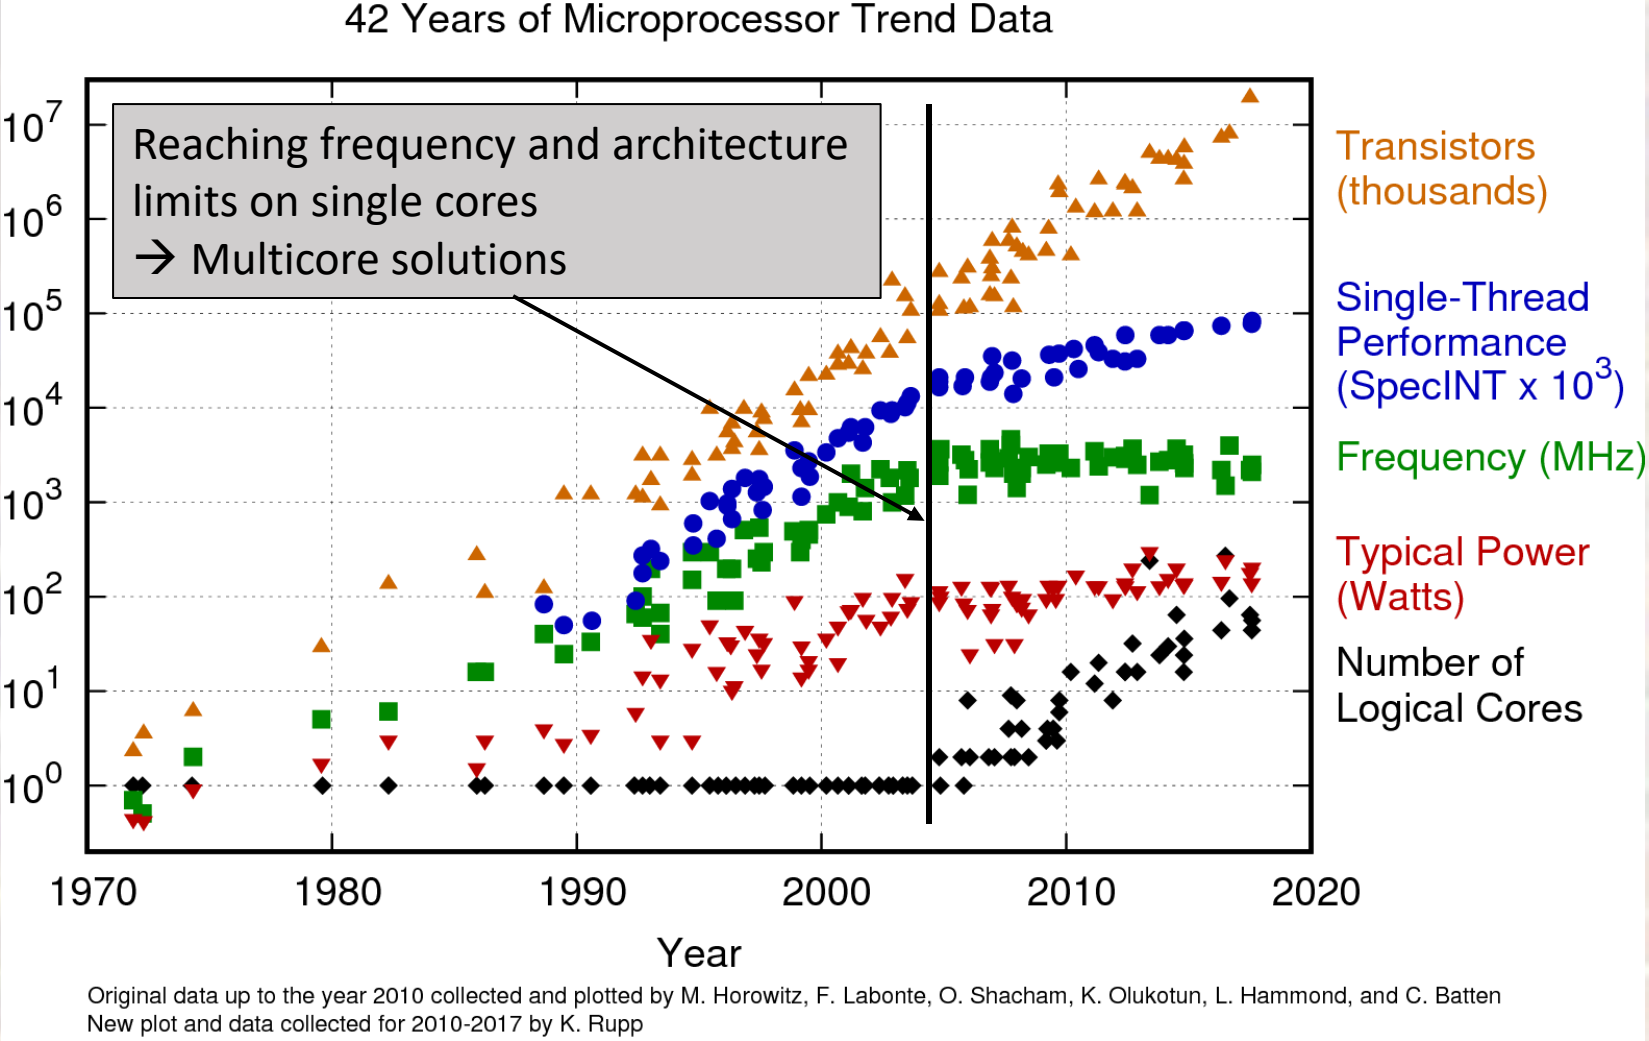

CMOS – Chip Level

- SPECint Performance – Including multi-core

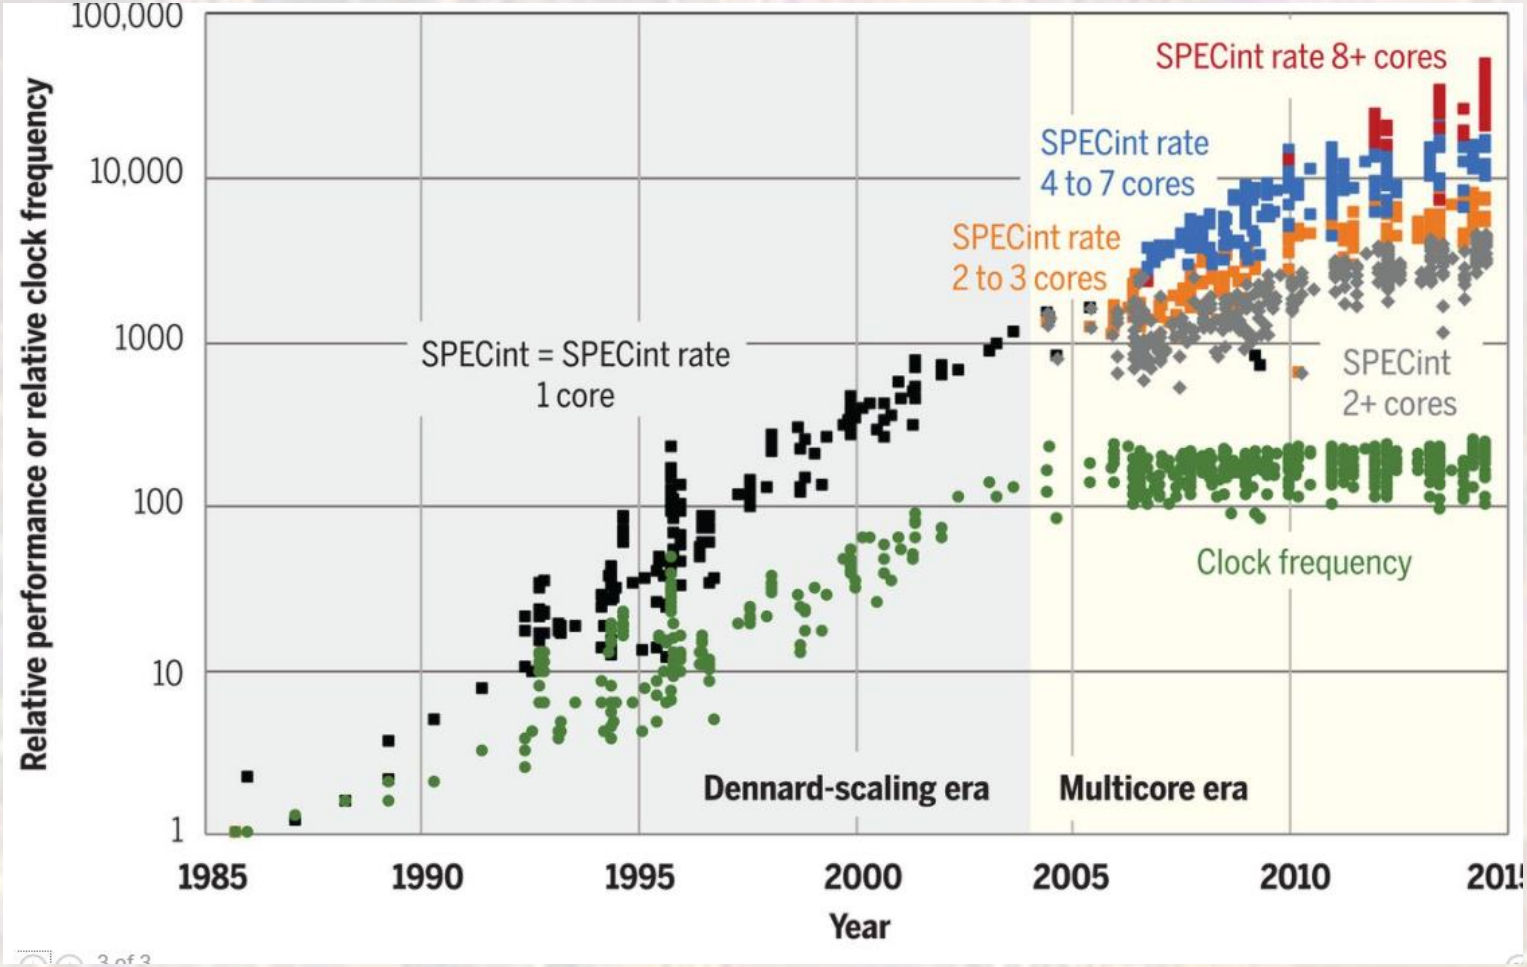

src: Science Magazine – June 2020

CMOS – Chip Level

- What about Cost $$$

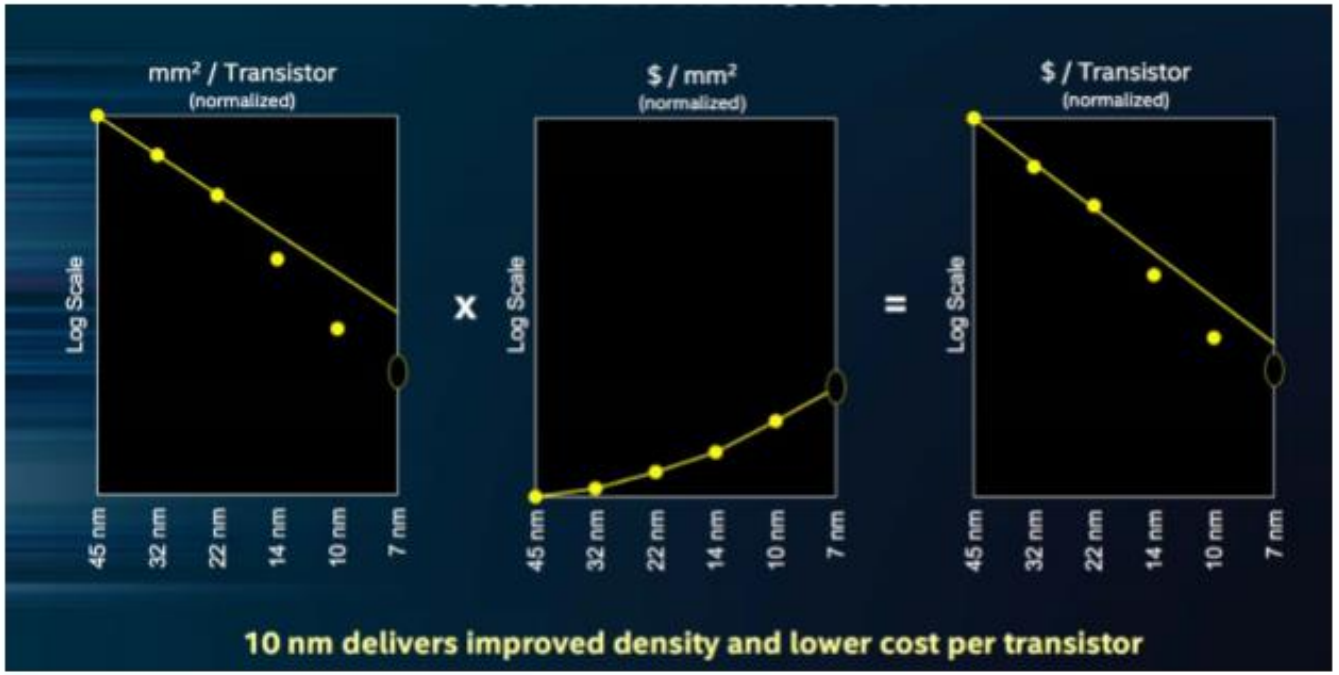

The minimum gate pitch of Intel's 10 nm process shrinks from 70 nm to 54 nm and the minimum metal pitch shrinks from 52 nm to 36 nm. These smaller dimensions enable a logic transistor density of 100.8 mega transistors per mm², which is 2.7x higher than Intel's previous 14 nm technology and is approximately 2x higher than other industry 10 nm technologies.

src: Intel

CMOS – Chip Level

- Design Processes
  - Tick-Tock
    - 2 year technology cycle
    - Tick - Existing processor design migrated to a new technology node
      - Higher speeds
      - Better power characteristics
    - Tock – New processor design introduced on an existing technology node
      - New features
      - Higher performance (IPC, …)

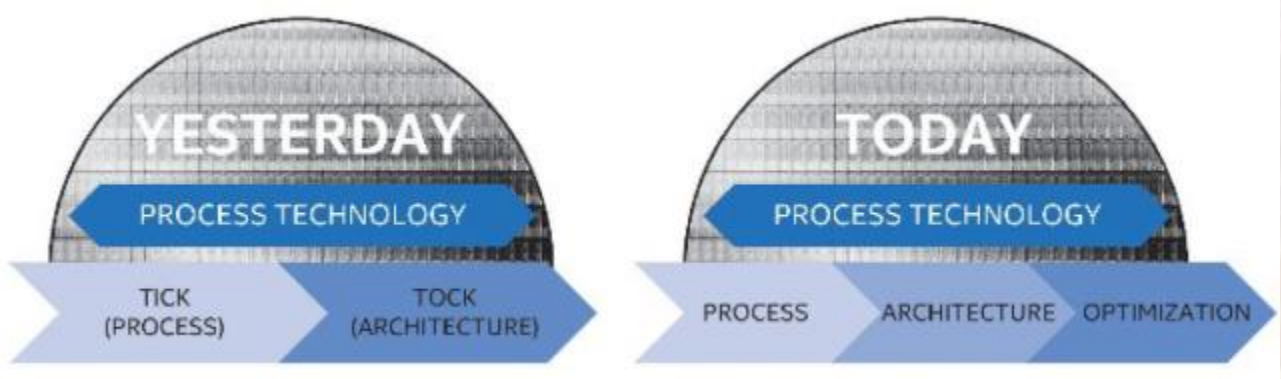

CMOS – Chip Level

- Design Processes
  - 3 Phase
    - 3 year technology cycle
    - Process
      - Slight modifications to existing architecture on a new technology node
    - Architecture
      - New architecture on the technology node
    - Optimization
      - Optimized architecture on the technology node