ΕΠΛ605
Προχωρημένη Αρχιτεκτονική Υπολογιστών
ιστοσελίδα θα ανακοινωθεί

Εισαγωγή:
Τάσεις Τεχνολογίας Υπολογιστών

Γιάννος Σαζεϊδης

Εαρινό Εξάμηνο 2017

READING

booksite.elsevier.com/9780123838728/sample_chapters/Hennessy_Chapter_1.pdf

Όλα τα Sections

Slides based on Elsevier Material

Computer Technology

- Amazing improvements:

  today can buy for $500 a mobile phone that has superior performance, more main memory and disk space than a $9million computer in 1975

- 9million/500 = 18000 times cost reduction
- Inflation: $1 1975 = $4.56 today
- 41million/500 = 80000 times cost reduction!

1975: Cray-1 (Wikipedia)

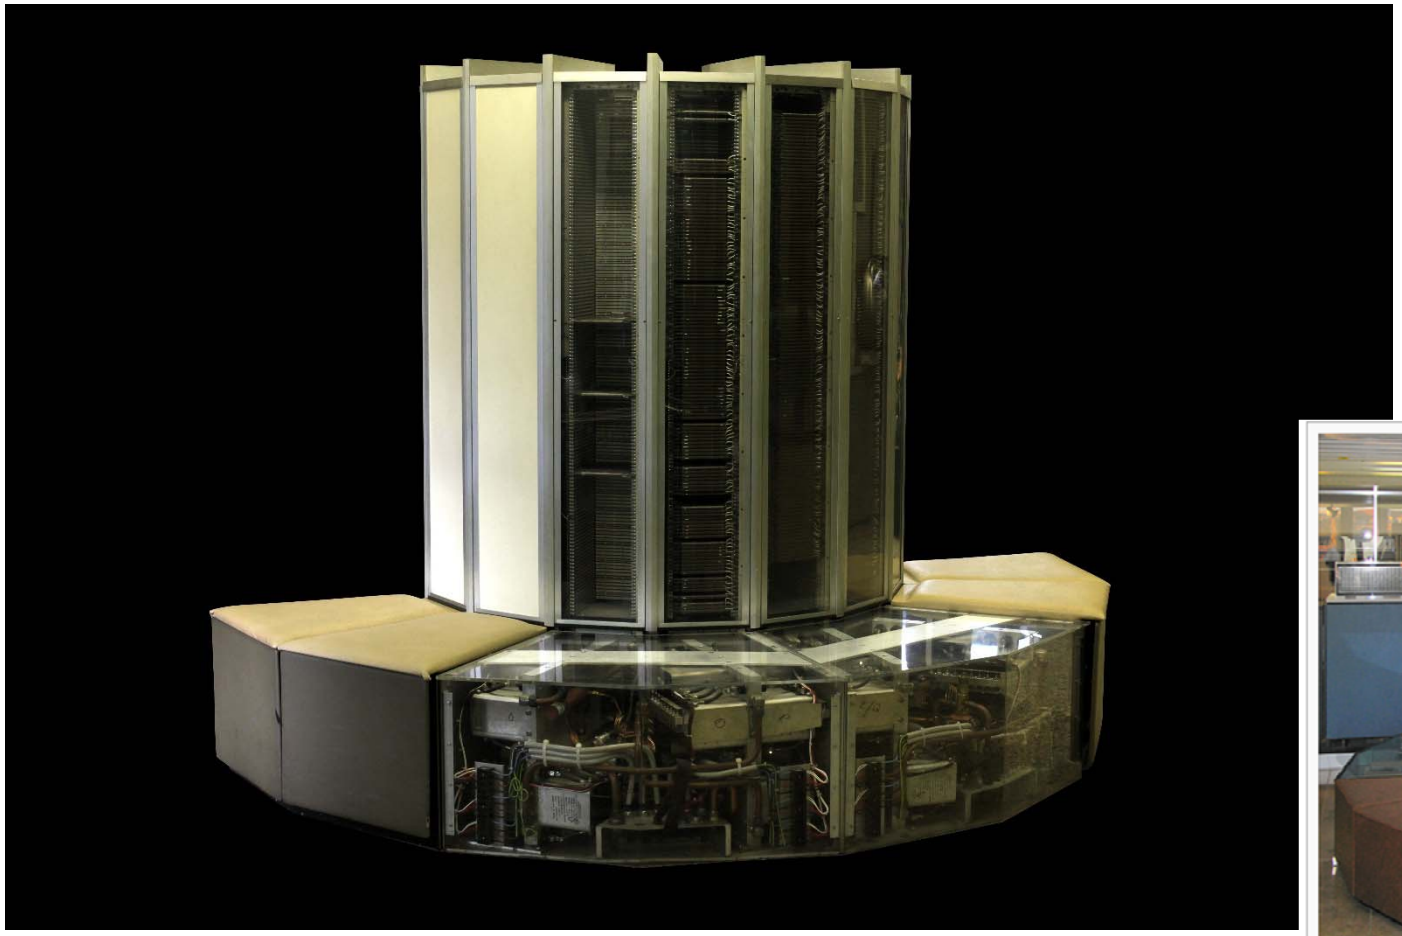

1975: Cray-1 (Wikipedia)

- Supercomputer
- Build by Cray Research
- Architect: Seymour Cray
- 64 bit machine: scalar and vector units
- Clock: 80MHz
- Technology: ECL, 2.5M transistors (1000s small chips)
- FLOPS: 160 MFLOPS
- Memory: 1M words (word: 64 bits, 8MB)
- Weight: 5.5 tons
- Cost: ~$9million
- Power: 115KW
- Sales Volume: ~80 units

2015: Galaxy S6 Edge

- Smartphone
- Build by Samsung
- 64 bit Exynos System on Chip
  - CPU: octa core, 4 A57, 4 A53
  - GPU: Mali 3D rendering engine
- Clock:
  - CPU: 2.1GHz, 1.5GHz
  - GPU: 600-700 MHz
- Technology 14nm FinFET, ~billion of transistors
- GPU FLOPS: >100s GFLOPS
- Memory: 3GB
- Storage: 32-128GB
- Weight: 132 g
- Cost: ~$500
- Battery: 2600 mAh
  - Battery Life depends on applications running
  - 10hrs @4V => 1W
- Sales Volume: 6million units in first month

Comparison

	1975	2015	Ratio
Transistors	2.5million	Billions	1000
Main Memory	1MB	3GB	3000
Clock	80MHz (12.5ns)	2.1GHz (0.48ns)	25
FLOPS	160M	100G	1000
Cost	$9million ($41mil)	$500	18000 (41000)
Weight	5.5 tons	132 g	~40000
Power	115KW	1W	115000
Sales	10s	millions	million

Computer Technology

- Amazing density/bandwidth improvements
  Source of improvements?

- Clock rate improves at a much slower rate

Bandwidth vs Latency

- Bandwidth or throughput – grows faster
  - Total work done in a given time
  - 10,000-25,000X improvement for processors
  - 300-1200X improvement for memory and disks
- Latency – grows slower
  - Time between start and completion of an event
  - 30-80X improvement for processors
  - 6-8X improvement for memory and disks

Bandwidth vs Latency

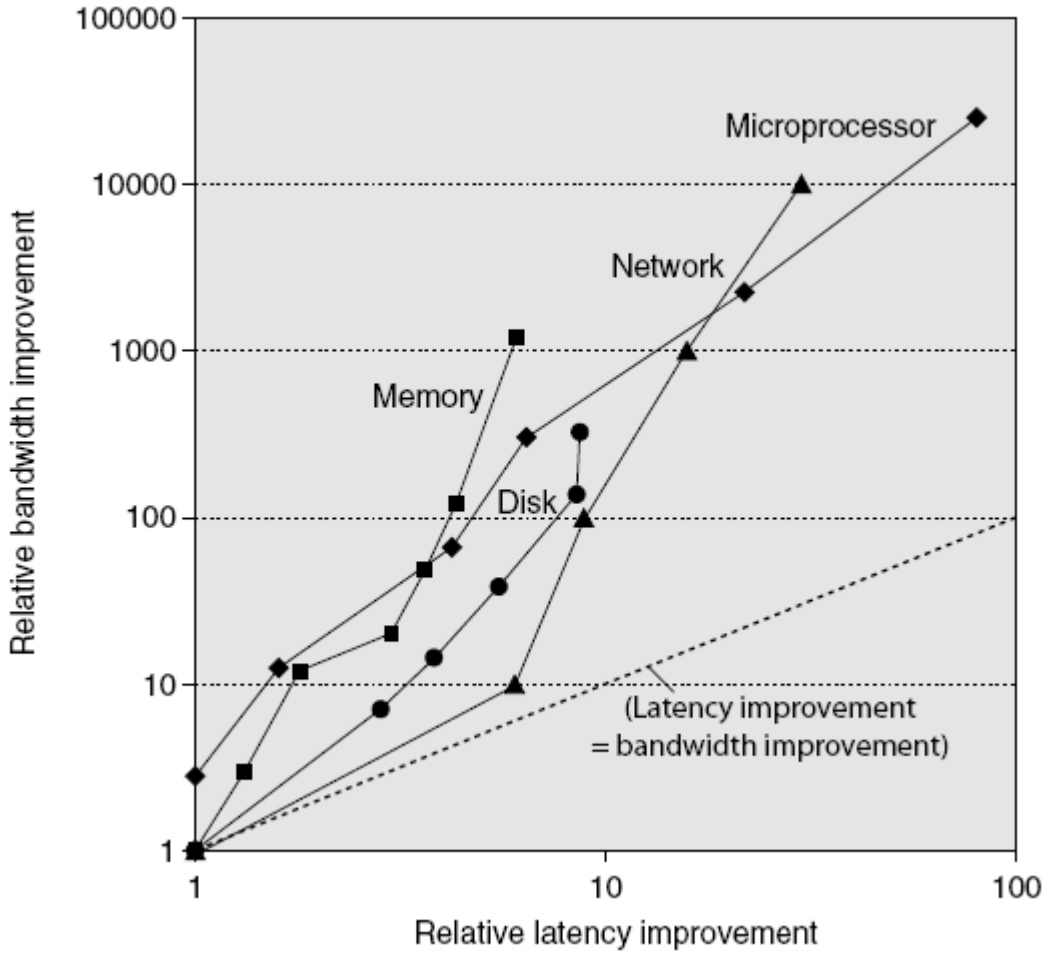

Log-log plot of bandwidth and latency milestones

Historic trend: bandwidth grows faster than latency improvement

Bandwidth vs Latency

- Τι είναι πιο εύκολο να καλύψω μια απόσταση Χ φορές πιο γρήγορα ή να την καλύψουν Χ παράλληλα

- Να επιταχύνω Χ φορές μια πράξη ή να εκτελώ Χ πράξεις παράλληλα

Computer Technology

- Source of improvements:
  - semiconductor technology building transistors
    - Feature size, energy per transistor
  - software
    - HLL compilers
    - Operating Systems
  - hardware design
    - Processor, main memory, storage, I/O, network
  - Cost per bit, operation decreases
  - Growth in sales volume
    - increasing number of applications and diverse market segments

Basic Economics

- Making feature size smaller is very expensive: multi-billion cost for a factory
- Such cost is affordable if it can be amortized in few years with billions of IC sales
- Larger volume helps lower non-recurring expenses (NRE) cost paid by each customer

E.g. $5billion new factory, 20million chips/year

NRE Cost/chip for 5 yrs = 5e9/(5x20e6) = $50

IC Sales

- 267 billion Integrated circuits expected to ship in 2016 (McClean Report)
- $203 billion sales (IC Insights)
  - 36.7 Billion Memory/Flash Units
  - 300 million pc/laptop processors (2015)
  - 23 million server processors – mostly Intel x86 (2015)
  - 14 billion ARM based chips (2015)

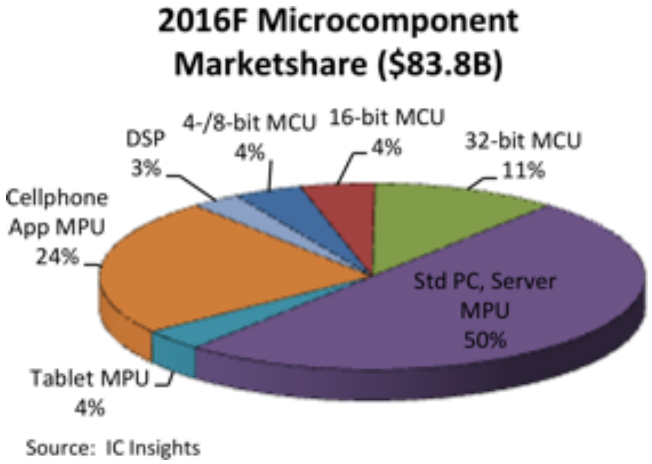

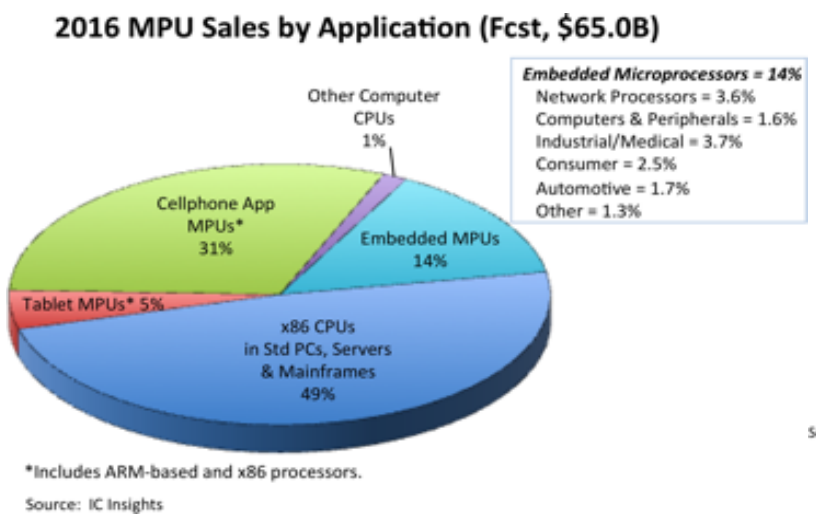

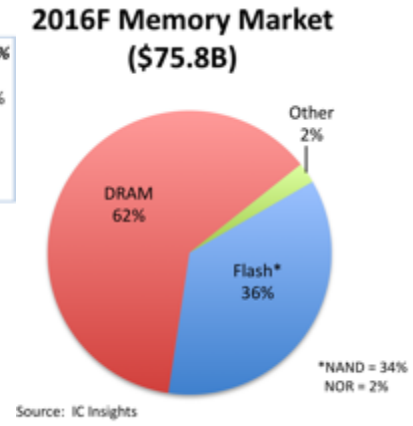

Classes of Computers

- Personal Mobile
  - e.g. smart phones, tablet computers
  - Emphasis on cost, video quality, energy efficiency and real-time
- Desktop/Laptop
  - Emphasis on cost-performance, graphics
- Servers
  - Emphasis on availability, scalability, throughput
- Clusters/Warehouse Scale Computers
  - Used for “Software/Infra/Platform as a Service (XaaS)”
  - Emphasis on availability and price-performance
  - Based on commodity parts
- Supercomputers
  - Emphasis: floating-point performance and fast internal networks
- Embedded:
  - Emphasis: price (and other depending on application/requirements, e.g. safety)

Embedded: Smart Card

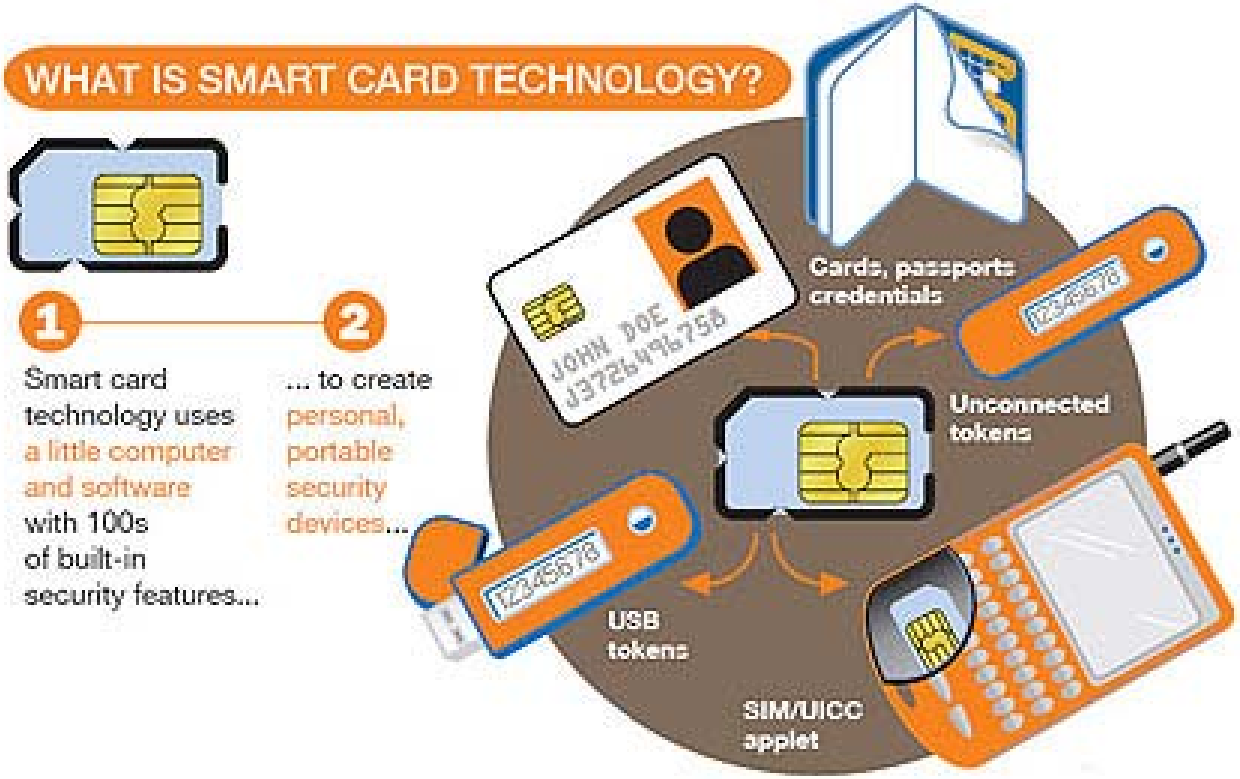

Personal Mobile: iPad

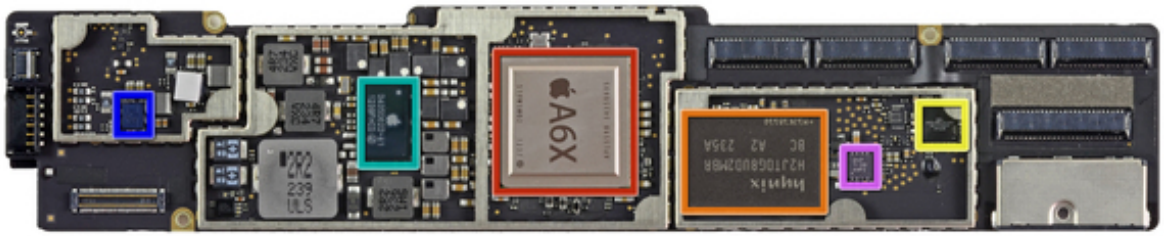

- Apple A6X Processor
- Hynix H2JTDG8UD2MBR 16 GB NAND Flash
- Apple 338S1116 Cirrus Logic Audio Codec
- 343S0622-A1 Dialog Semi PMIC
- Apple 338S1077 Cirrus Logic Class D Amplifier
- QVP TI 261 A9P2

2 x 4Gb Elpida LP DDR2

Cloud Server (x86 based)

Racks with Servers: Data Center

Farm with 10000s servers
Interconnect architecture very important

Top Supercomputer: TaihuLight (June 2016)

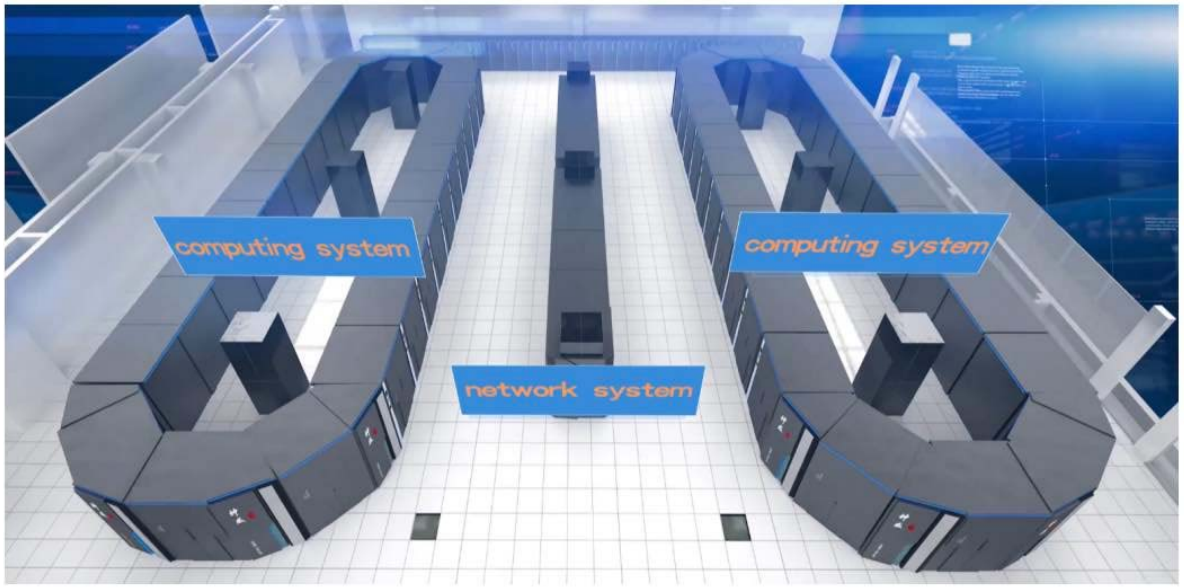

Figure 4: Overview of the Sunway TaihuLight System

Architecture

The Sunway TaihuLight uses a total of 40,960 Chinese-designed SW26010 manycore 64-bit RISC processors based on the Sunway architecture.[5] Each processor chip contains 256 processing cores, and an additional four auxiliary cores for system management that are also RISC cores just more fully featured, for a total of 10,649,600 CPU cores across the entire system

Sunway TaihuLight	
Active	June 2016
Location	National Supercomputer Center, Jiangsu, China
Architecture	Sunway
Power	15 MW
Operating system	Sunway RaiseOS 2.0.5 (based on Linux)
Memory	1.31 PB
Speed	93 PFLOPS
Cost	1.8 billion Yuan (US$273 million)
Purpose	Oil prospecting, life sciences, weather forecast, industrial design, drug research
Web site	http://demo.wxmax.cn/wxc/index.php

Key Driving Applications

- **Scientific: weather prediction, genome sequencing**
  - Need: large memory, heavy-duty floating point, availability (hw)
  - Intel, AMD, Nvidia based, custom
- **Cloud services and Social Media: amazon, google, Microsoft facebook, liknkedin,**
  - Need: availability (hw+sw), quality of service
  - Intel based, fpga
- **Commercial: database/web serving,**
  - Need: data movement, high memory + I/O bandwidth
  - Intel based, Oracle
- **Home office, multimedia, games**
  - Need: integer, memory bandwidth, graphics
  - Intel/AMD/NVIDIA based
- **Mobile apps:**
  - Need: low power, integer performance, integrated graphics
  - ARM based by Samsung, Qualcomm;
- **Embedded Applications**
  - Smart Cars (functional safety): Autonomous driving
  - Internet of Things: RFID on everything, temperature, air, water monitoring, traffic, smart houses, survailance(security)
  - Deeply Embedded: disposable “smart dust” sensors
  - Need: low power, low cost (not always) – e.g. safety applications
  - ARM based chips, dedicated digital signal processing
  - Large collection of embedded processors needs backing up from servers to do processing, aggregation, coordination etc

Βασική Τεχνολογία Υλοποίησης

- Processor logic and memory arrays implemented with silicon based transistors
- DRAM with transistors and capacitors
- Metal wires for connectivity
- Feature Size:….,0.045u,*0.032u*,0.028u,0.022u,*0.014u*,*0.010u*…
- Συνέπεια σμίκρυνσης στον ίδιο χώρο περισσότερους πόρους (transistors, functional units, memory cells)

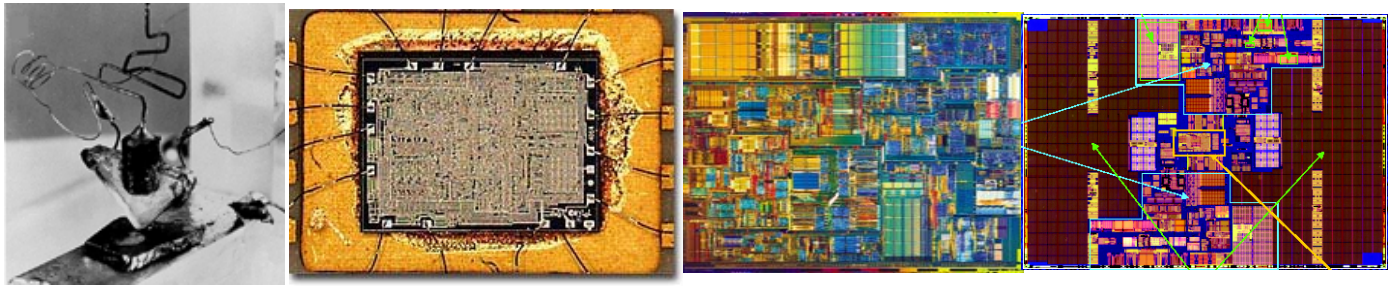

1947-10^0 1971-10^3 2002-10^7 2005-10^9

Πώς αξιοποιώ τα transistors?

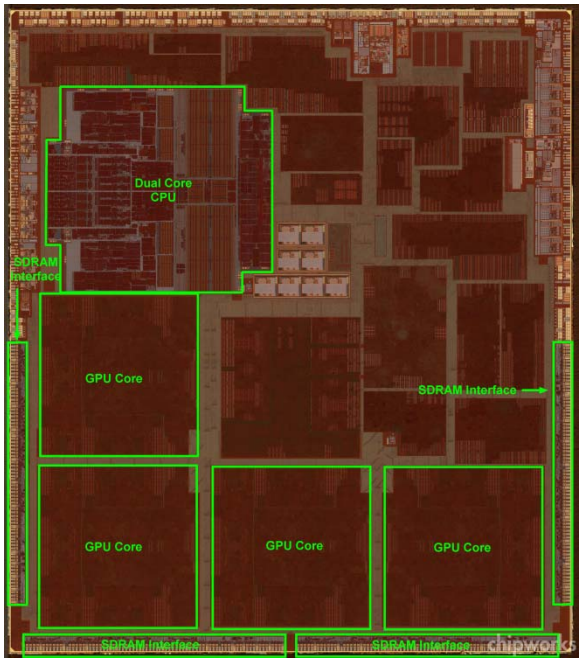

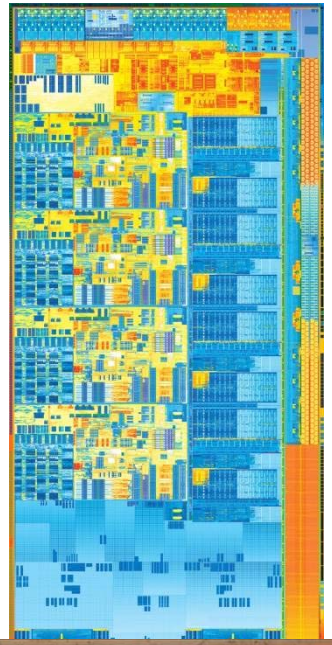

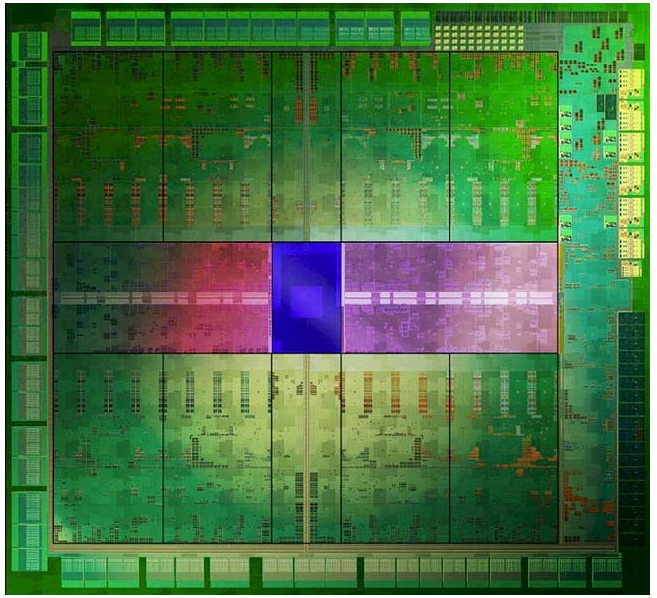

https://www.youtube.com/watch?v=aCOyq4YzBtY

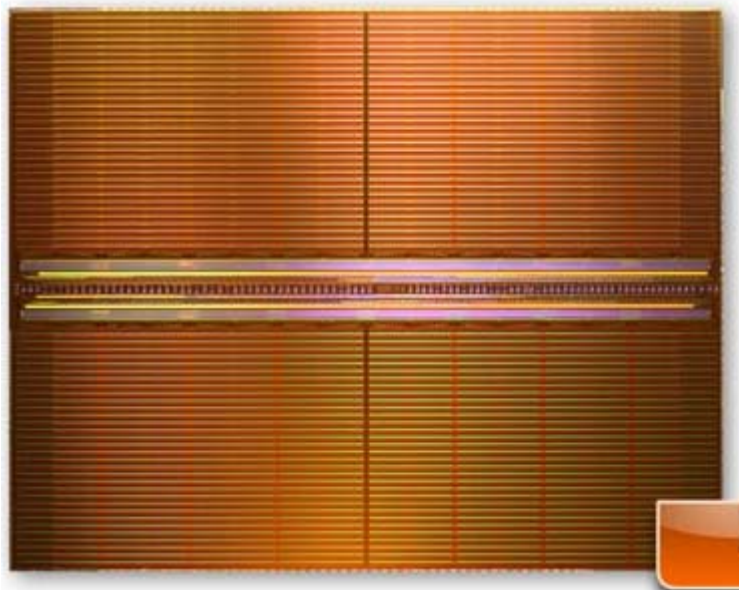

Trends in Technology: Growth

- Integrated circuit technology
  - Transistor density: 35%/year (Moore's law)
  - Die size: 10-20%/year
  - Integration overall: 40-55%/year
- DRAM capacity: 25-40%/year
- Flash capacity: 50-60%/year
  - 15-20X cheaper/bit than DRAM
- Magnetic disk technology: 40%/year
  - 15-25X cheaper/bit than Flash
  - 300-500X cheaper/bit than DRAM
- Network Technology (routers and links)

Amazing Density Trends

Clock, Energy, Power

- Transistors keep getting smaller and faster
- Moore's "Law": doubling transistor density every 24-36 months
- Transistors burn energy
  - if 2x transistors in next generation chip
  - for same energy cut energy/transistor by 0.5x
- Challenge: energy per transistor is scaling slower with smaller feature size
  - Supply voltage is not scaling (Dennard Scaling)

Technology scaling: Challenges

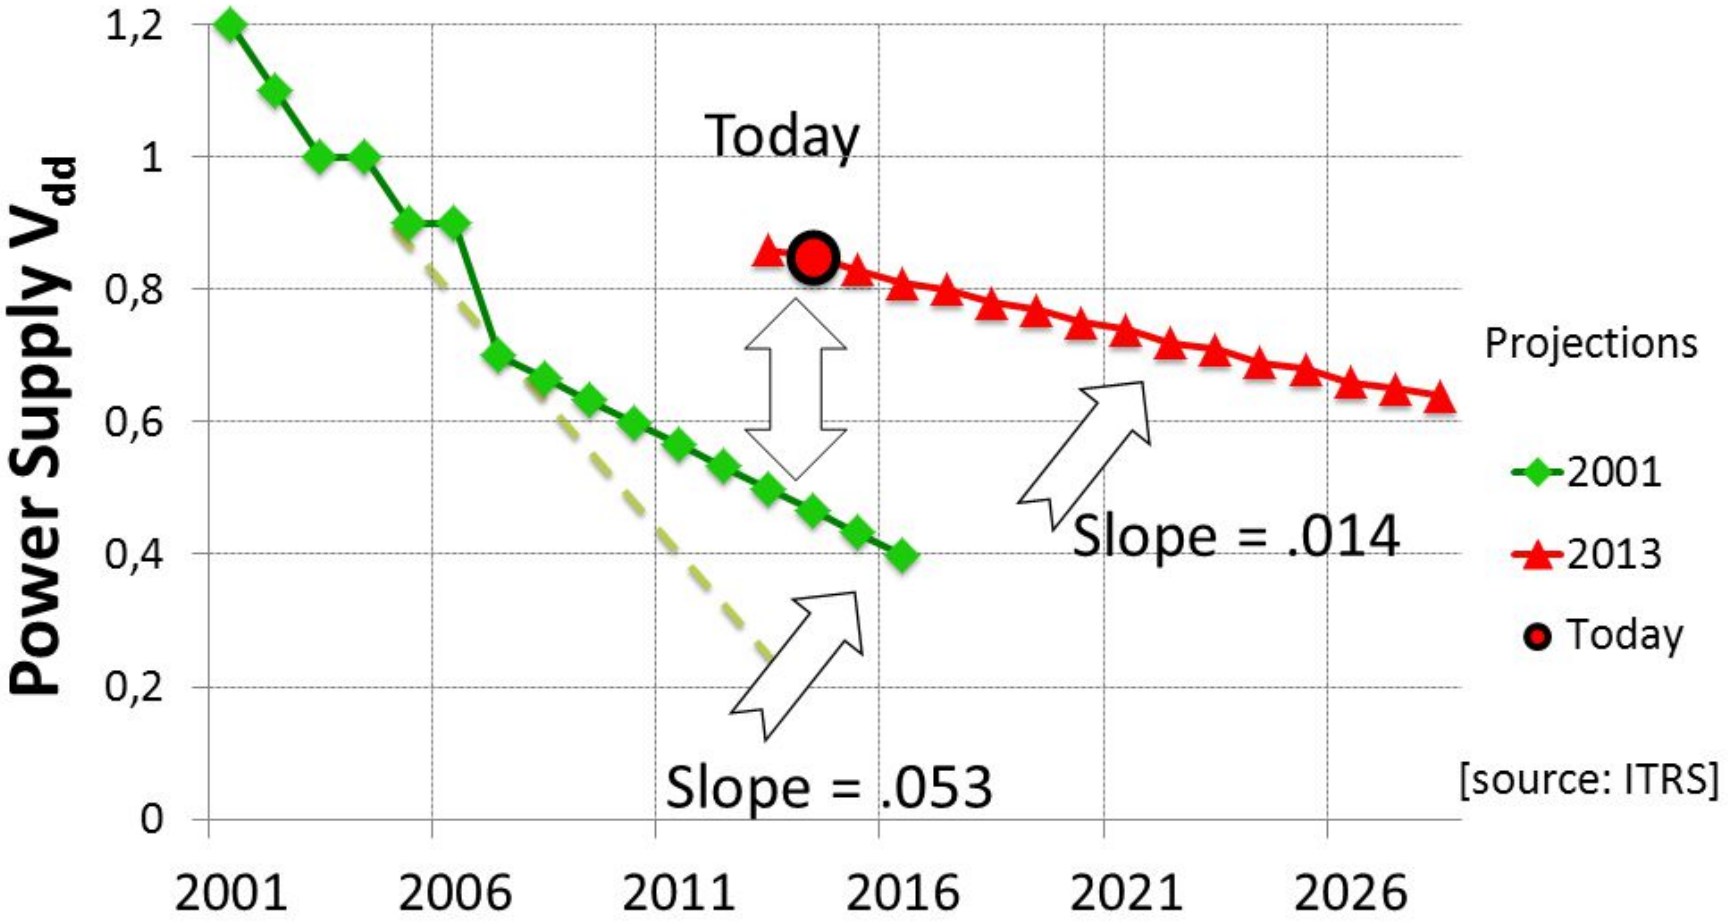

Clock, Energy, Power

- Power: rate of burning energy Energy/Time
  - Depends on clock rate
- Higher power => higher temperature
- All platforms have power envelopes (fixed ☹)
- Trade-off:
  - lower clock rate
  - For same power create room to burn more energy
  - More transistors used to do more work (How to use them?)

- Processor clock speed not scaling in last decade
- Emergence of multicores, GPU and other accelerators

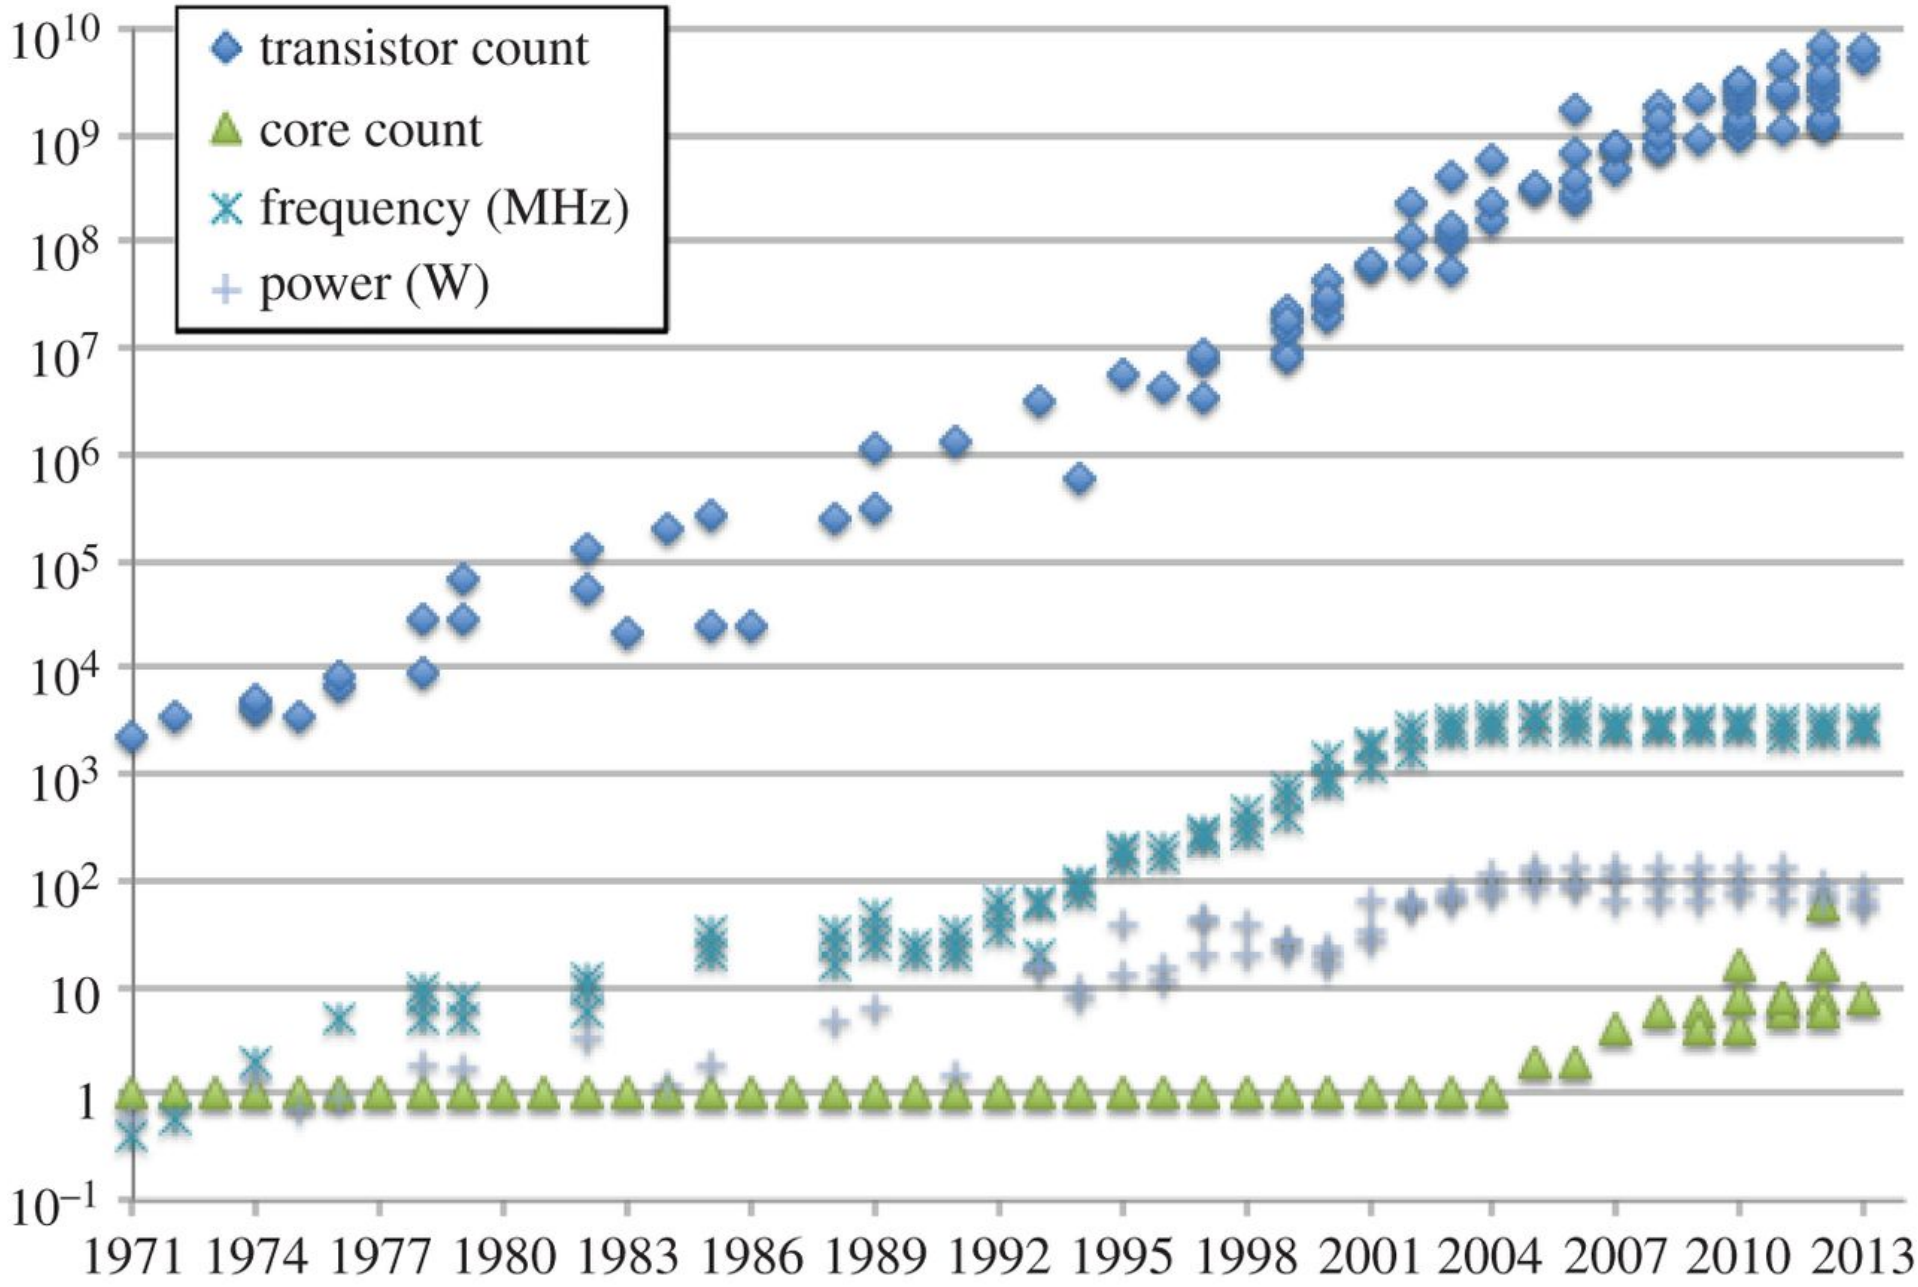
transistor count
core count
frequency (MHz)
power (W)
10^{10}
10^{9}
10^{8}
10^{7}
10^{6}
10^{5}
10^{4}
10^{3}
10^{2}
10
1
10^{-1}
1971 1974 1977 1980 1983 1986 1989 1992 1995 1998 2001 2004 2007 2010 2013

Single Processor Performance

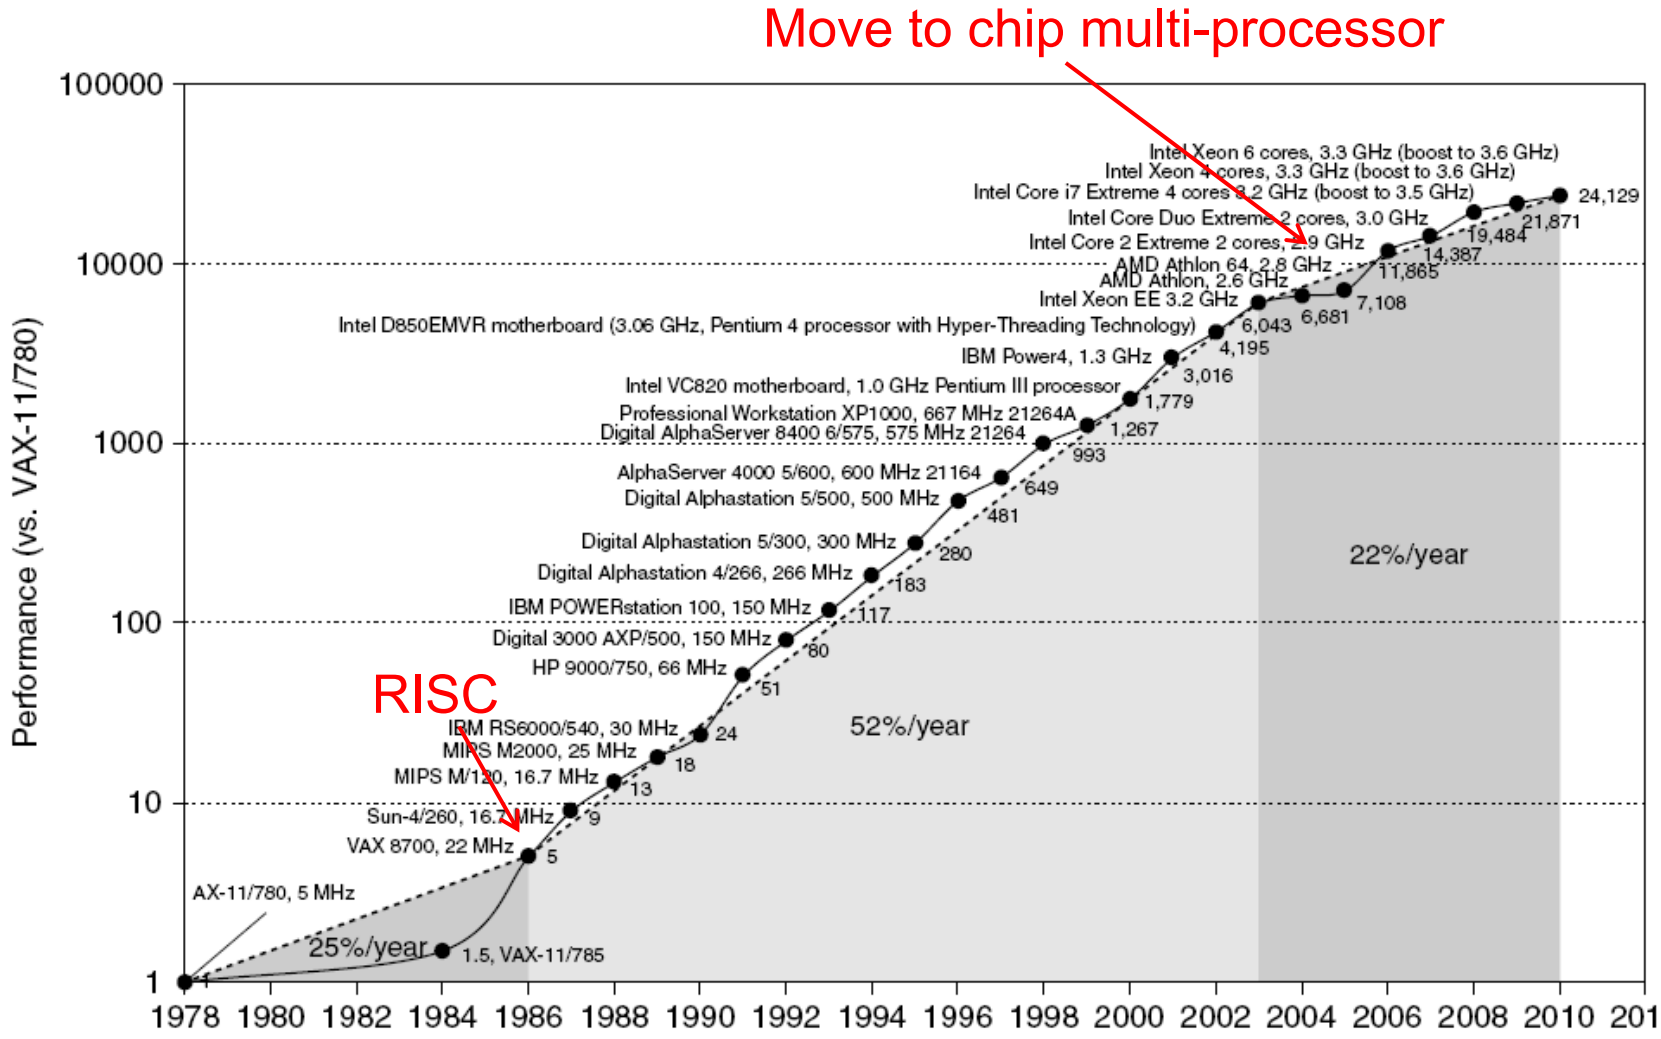

RISC Reduced Instruction Set Computer

Transistors and Wires

- Feature size
  - Minimum size of transistor or wire in x or y dimension
  - 10 microns in 1971 to .010 microns (10nm) in 2016 (10/0.01~ 1000x reduction in 1D and ~1000000x reduction in 2D)
  - Scaling of speed of transistors vs wires not-uniform (quadratic vs linea

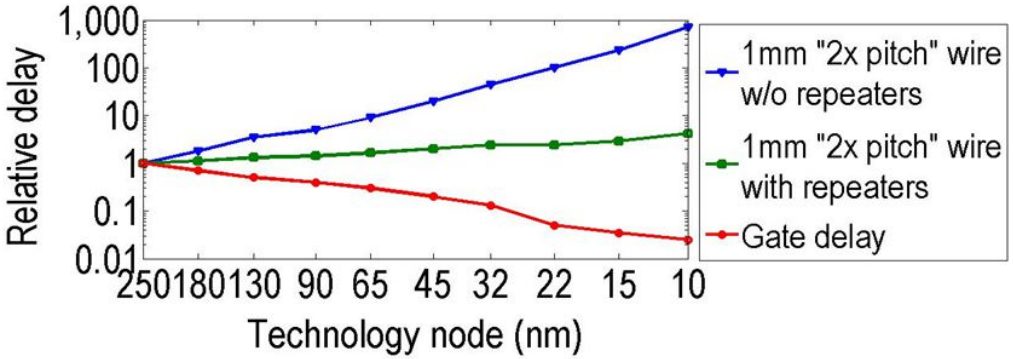

- ***Non-uniform scaling supports localized activity, caching, and communication***

Μέλλον του Moore's Law

- Roadmaps Until ~2025 (10nm, 7nm, 5nm)
  - Will continue (prototypes in lab)
  - Costs will increase (need larger volumes)
- Long-Term (>~2025 years)
  - 3D stacking
  - May need new technology (e.g. quantum)
  - DNA based
  - We can do better (e.g., human brain, approximate computing)

Trends in Computer Architecture

Architects Role: decide how to use the transistors, what frequency to run the processor/memory UNDER power, thermal, reliability, area, cost CONSTRAINTS maximize profit

PAST:

- For several years improve performance with more Instruction Level Parallelism (ILP)
  - Pipelining
  - Out-of-order execution
  - Superscalar execution
  - Prediction and Speculative execution
  - Clock Rate Increase

- Improvements in Memory Hierarchy: capacity and access time
  - Multi-level Caches,
  - Replacement Policies
  - Prefetching

- Multiprogram parallelism: running independent programs on different processors

Trends in Computer Architecture

PRESENT:

- Continue to improve ILP and memory hierarchy
  - Single processor performance improvement key business requirement
  - technology and program properties make it challenging
  - Amdahl's Law makes it essential
- And multiprogram parallelism

- Other types of parallelism pursued
  - Data-level parallelism (DLP)
  - Thread-level parallelism (TLP)
    - Require parallel architectures and explicit restructuring of the applications (e.g. parallel programming 372)
  - Request Level Parallelism (RLP) (Cloud services)
- Also accelerators – e.g. deep neural nets for classification

Current Parallelism in Computer Classes (processors/cores)

- Personal Mobile Device (PMD) 1 <10
- Desktop Computing 1-2 10-100
- Servers 1-10s 10-1k
- Clusters/Warehouse Scale Computers 1K to 100K (x10)
- Supercomputers 40k 1e6
- Embedded Computers can vary

Parallelism

- Classes of architectural parallelism:
  - Instruction-Level Parallelism (ILP)
  - Multiprogram (multicores,multiprocessors)
  - Vector architectures/Graphic Processor Units (GPUs)
  - Thread-Level Parallelism (multicores,multiprocessors)
  - Request-Level Parallelism (multicores,multiprocessors)

PP is not New, March 1961:

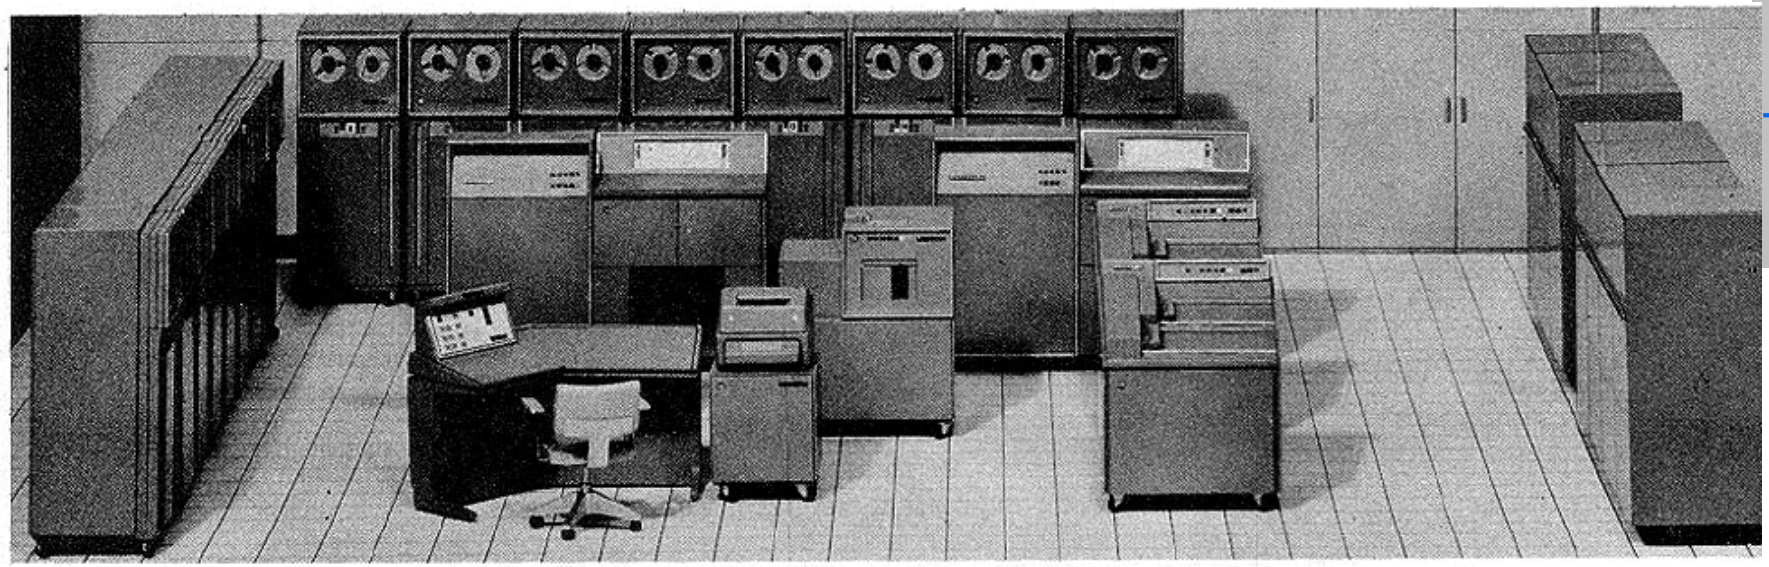

Finally hatched

BURROUGHS ANNOUNCES THE B5000

IT WAS A LONG TIME COMING – but the Burroughs Corporation B5000, that firm's first entry in the solid state computer field, promises to be an interesting addition to the industry.

Featuring an add time of three microseconds and a six microsecond memory cycle, a maximum system may include up to eight memory modules, each containing 4,096 49 bit words.

An interesting feature of the machine lies in the fact that two central processors may be employed simultaneously for what Burroughs calls "true parallel processing."

B5000 will rent from $13,500 to $50,000 a month. The sale price range is $540,000 to $2,000,000. A more detailed article on the B5000 will appear in an early issue of DATAMATION.

Parallel Processing is not New

- Multiprocessors have been around for decades
- Using specialized or commodity parts
- Parallel programming for certain domains very well developed art
- Many parallel programming languages and frameworks
- Numerous efforts for automatic (compiler) parallelization and vectorization

What is new

- Technology trends
- Embrace of multi/many cores across industry
- Emergence of large scale RLP
  - More throughput
- No silver bullet (perfect solution) for all applications
- Specialization and heterogeneity
  - Small cores, big cores, gpus, accelerators, neural nets etc

Αρχιτεκτονική Υπολογιστών

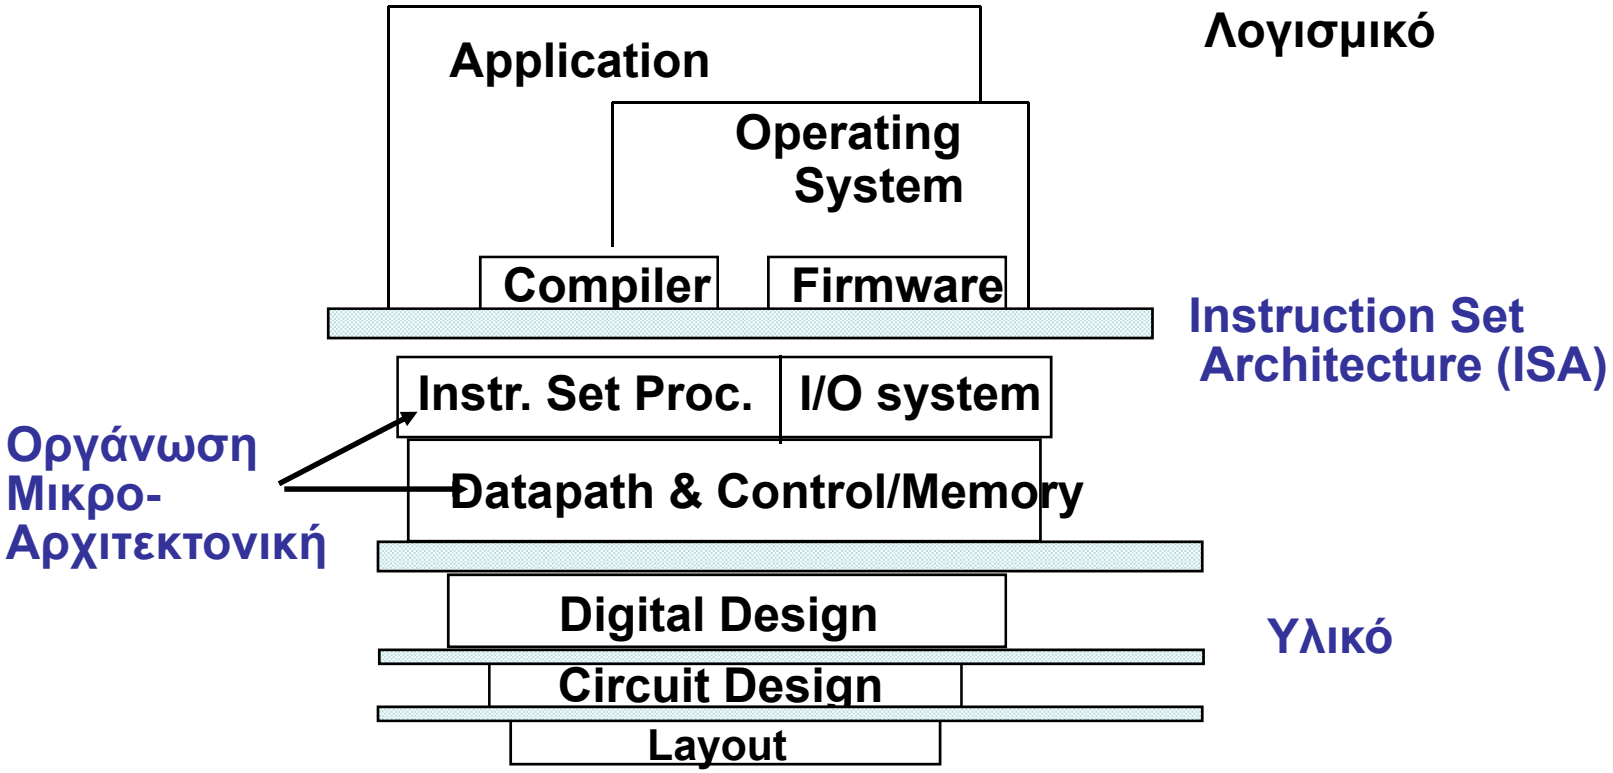

- Η ΑΣΕ (ISA) διασύνδεση μεταξύ λογισμικού και υλικού
- Η οργάνωση περιγραφή ροής και αποθήκευσης πληροφοριών και έλεγχου ξεχωριστά από την τεχνολογία υλοποίησης

ΑΣΕ (ISA)

Name	Number	Use	Preserved across a call?
$zero	0	The constant value 0	N.A.
$at	1	Assembler temporary	No
$v0–$v1	2–3	Values for function results and expression evaluation	No
$a0–$a3	4–7	Arguments	No
$t0–$t7	8–15	Temporaries	No
$s0–$s7	16–23	Saved temporaries	Yes
$t8–$t9	24–25	Temporaries	No
$k0–$k1	26–27	Reserved for OS kernel	No
$gp	28	Global pointer	Yes
$sp	29	Stack pointer	Yes
$fp	30	Frame pointer	Yes
$ra	31	Return address	Yes

PC: program counter points to next instruction to execute

Instruction type/opcode
Data transfers
LB, LBU, SB
LH, LHU, SH
LW, LWU, SW
LD, SD
L.S, L.D, S.S, S.D
MFC0, MTC0
MOV.S, MOV.D
MFC1, MTC1
Arithmetic/logical
DADD, DADDI, DADDU, DADDIU
DSUB, DSUBU
DMUL, DMULU, DDIV, DDIVU, MADD
AND, ANDI
OR, ORI, XOR, XORI
LUI
DSLL, DSRL, DSRA, DSLLV, DSRLV, DSRAV
SLT, SLTI, SLTU, SLTIU
Control
BEQZ, BNEZ
BEQ, BNE
BC1T, BC1F
MOVN, MOVZ
J, JR
JAL, JALR
TRAP
ERET
Floating point
ADD.D, ADD.S, ADD.PS
SUB.D, SUB.S, SUB.PS
MUL.D, MUL.S, MUL.PS
MADD.D, MADD.S, MADD.PS
DIV.D, DIV.S, DIV.PS
CVT._._
C.__.D, C.__.S

ΑΣΕ (ISA)

Οργάνωση

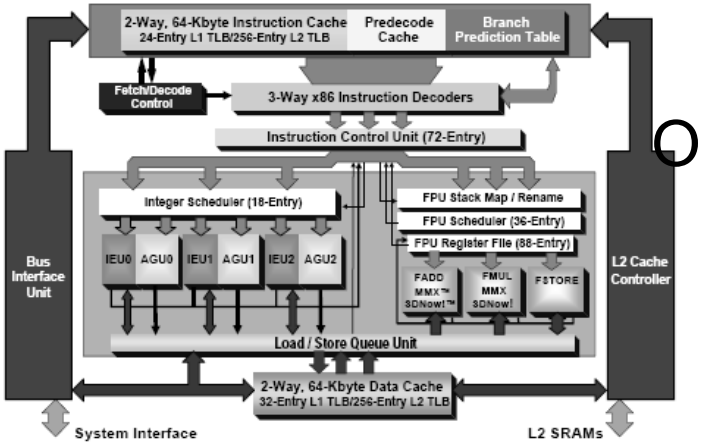

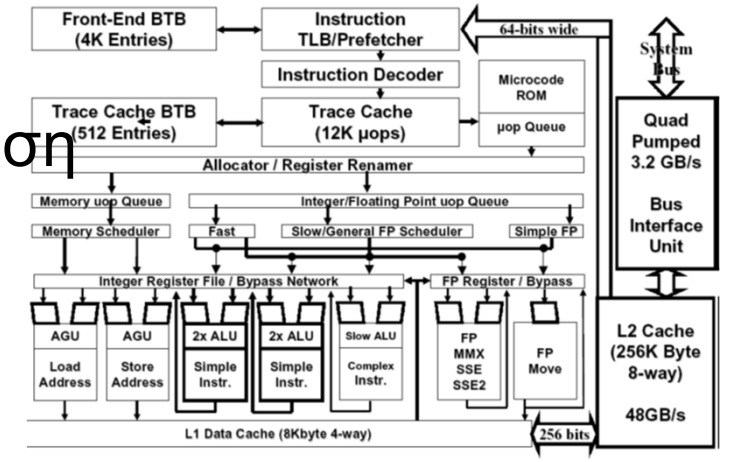

Υλικό/Hardware

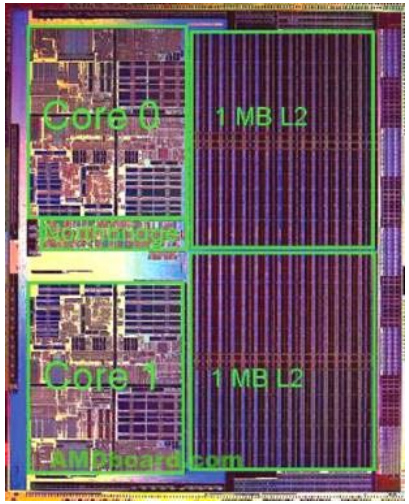

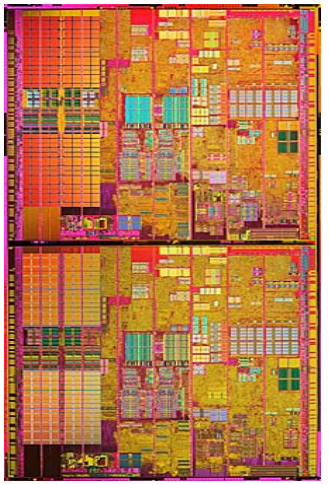

Defining Computer Architecture

- Instruction Set Architecture (ISA) design
  - i.e. decisions regarding:
    - registers, memory addressing, addressing modes, instruction operands, available operations, control flow instructions, instruction encoding
  - example ISAs: x86, ARM, MIPS, IBM
    - 32 bit and 64 bit
    - CISC, RISC
  - Extending an ISA: new operations, e.g. bit manipulation to facilitate security, new data types e.g. 256-bit operands

Defining Computer Architecture

- Computer organization:
  - Hardware implementation of the ISA
    - Number of pipe stages, size of caches, number of functional units, prediction, prefetching…
  - Design to maximize performance within constraints: cost, power, and availability
  - Many options for implementing the ISA (huge and challenging design space)
    - Critical making right choices
  - Hardware configurable and dynamic: challenging how to configure it and use it efficiently

Γιατί να πάρετε μάθημα αρχιτεκτονικής υπολογιστών

- Κατανόηση της οργάνωσης, των περιορισμών και δυνατοτήτων των μοντέρνων και μελλοντικών υπολογιστών
- Σχεδιασμός σε ψηλό επίπεδό - όχι λεπτομερής σχεδιασμός
- Καλύτερη γνώση πώς επηρεάζει υλικό την επίδοση των προγραμμάτων και πως γράφουμε καλύτερα προγράμματα
  - Τι περιορίζει/επηρεάζει την επίδοση;
  - Πως βελτιώνουμε την επίδοση;
- Εξοικείωσης με εργαλεία και τεχνικών για ανάλυση της επίδοσης ενός προγράμματος.

Πληροφορίες για το ΕΠΛ605

- Διδάσκων: Γιάννος Σαζεϊδης
- Διαλέξεις: 4.30-6 Δευτέρα- 4.30-6 Πέμπτη
- Φροντιστήριο Πέμπτη 6-7
- Εργαστήριο: Δευτέρα 10.30-12.30
- Βοηθός: Πέτρος Παναγή
- Ιστοσελίδα: www.cs.ucy.ac.cy/courses/EPL605

- Τι θα μάθετε/στόχοι:
  - Τεχνολογία Υπολογιστών – Τάσεις, ευκαιρίες και περιορισμοί
  - Μεθοδολογία αξιολόγησης και σύγκρισης απόδοσης επεξεργαστών
  - Βασικές αρχές: pipelining και memory hierarchy
  - Προχωρημένες Έννοιες: Παραλληλισμός (ILP, DLP, TLP)
  - Χρήση διαφορών εργαλείων για εξερεύνηση σε βάθος κάποιων πτυχών της αρχιτεκτονικής

Πληροφορίες για το ΕΠΛ605

- Προαπαιτούμενα
  - Οργάνωση Υπολογιστών και Συμβολικός Προγραμματισμός (ΕΠΛ221)
    - Αρχιτεκτονική συνόλου εντολών (ISA): πχ MIPS, x86 κτλ
    - Διάδρομος δεδομένων (datapath)
    - Μονάδα ελέγχου (control unit)
    - Σχεδιασμός ενός ψηφιακού συστήματος/απλού επεξεργαστή
    - Τα βασικά σε σχέση με κρυφή μνήμη (caches) και διασωλήνωση (pipelining)
    - Διαχείριση Μνήμης (εικονίκη/πραγματική μνήμη, caches)
  - Καλή Γνώση C και UNIX
- Βιβλιογραφία: Computer Architecture a Quantitative Approach, Henessy & Patterson 5th edition, άρθρα, web

Πληροφορίες για ΕΠΛ605

- Εργαστήριο:
  - εισαγωγή εργαλείων (πχ scripting, simplescalar,mars, gprof, pin,perf)
  - αξιολόγηση εργασιών
  - συζητήσεις, παρουσιάσεις
  - Χρειάζεστε UNIX λογαριασμό
- Αξιολόγηση
  - Εργασίες 25-30 %
  - Περίληψη/παρουσίαση άρθρων 5%
  - Συμμετοχή 5%
  - Τελική Εργασία (project) 15% (μεταπτυχιακούς)
  - Τελική – 45-55%