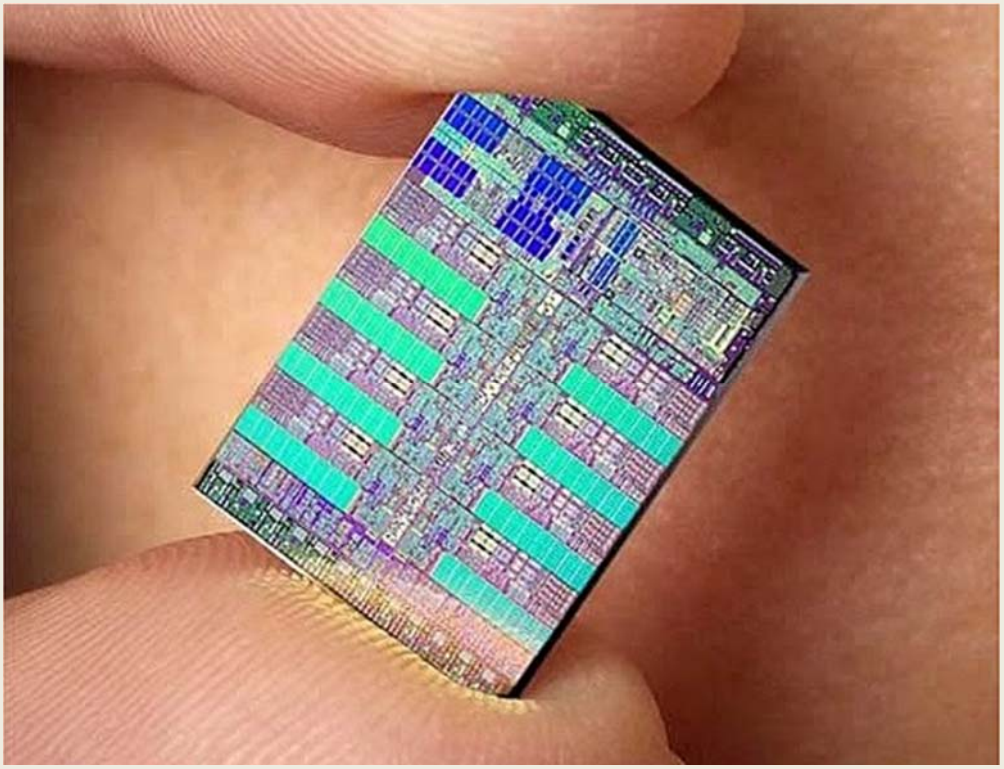

Microprocessor Technology: Past, Present, and Future

Professor Won Woo Ro, School of Electrical and Electronic Engineering
Yonsei University

The 1st Microprocessor

The Intel® 4004 microprocessor, introduced in **November 1971**

An electronics revolution that changed our world.

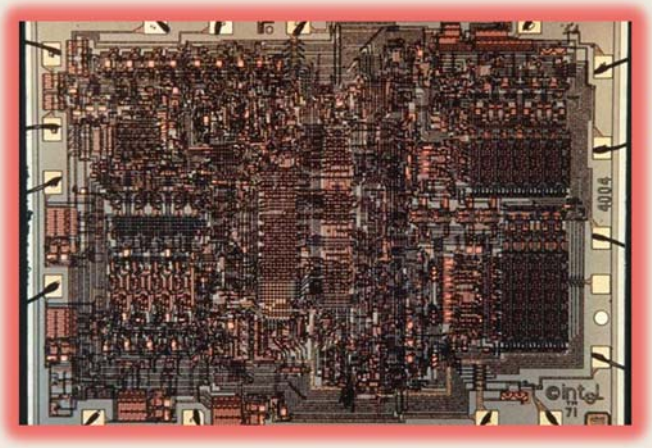

There were no **customer-programmable microprocessors** on the market before the 4004.

It propelled **software** into the limelight as a key player in the world of digital electronics design.

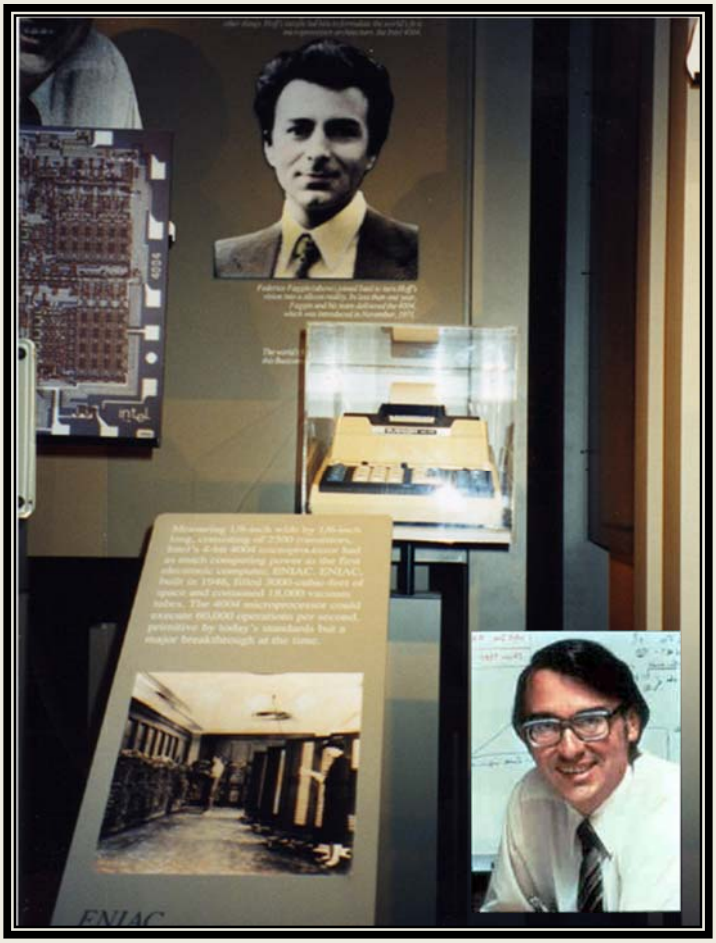

4004 Microprocessor Display at New Intel Museum

A Japanese calculator maker (Busicom) asked to design:

A set of **12 custom logic chips** for a line of **programmable calculators**.

Marcian E. "Ted" Hoff

Recognized the integrated circuit technology (of the day) had advanced enough **to build a single chip, general purpose computer**.

Federico Faggin

to turn Hoff's vision into a silicon reality. (In less than one year, Faggin and his team delivered the 4004, which was introduced in November, 1971.)

The world's first microprocessor application was this Busicom calculator. (sold about 100,000 calculators.)

Measuring **1/8 inch wide by 1/6 inch long**, consisting of 2,300 transistors, Intel's 4004 microprocessor had as much computing power **as the first electronic computer, ENIAC**.

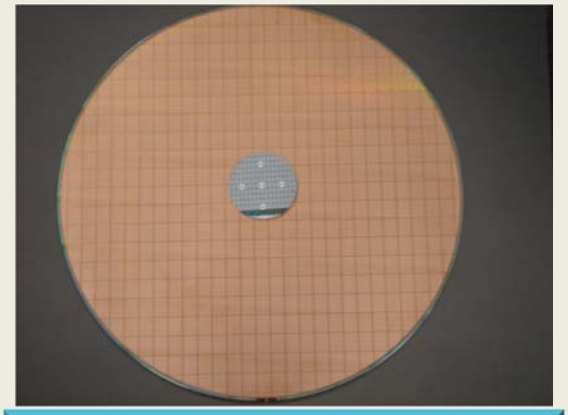

2 inch 4004 and 12 inch Core™2 Duo wafer

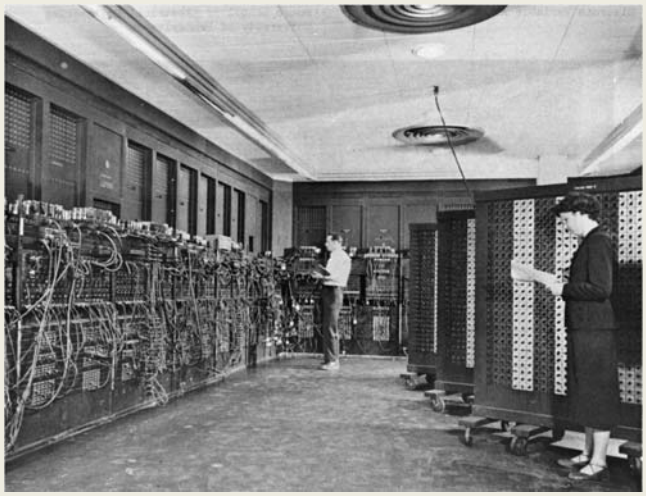

ENIAC, built in 1946, filled 3000-cubic-feet of space and contained 18,000 vacuum tubes.

The 4004 microprocessor could execute 60,000 operations per second

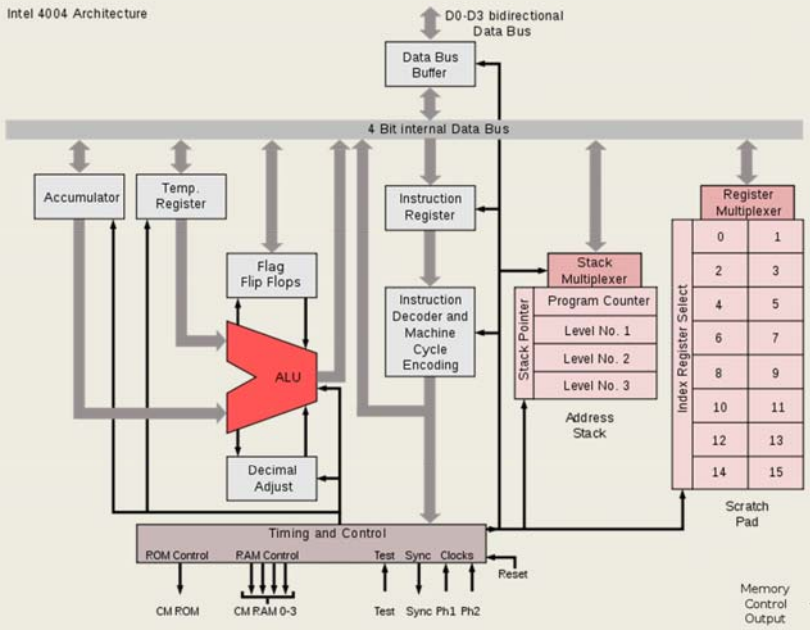

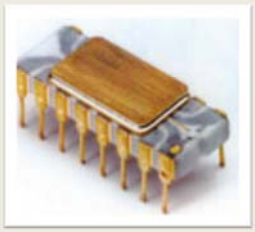

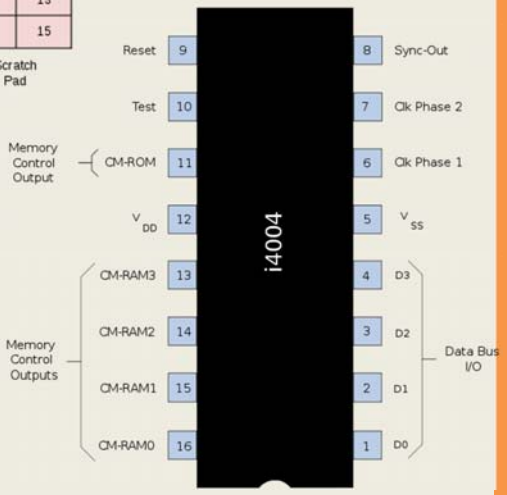

Running frequency: 108 KHz

Founders wanted to name their new company *Moore Noyce*. However the name sounds very much similar to "*more noise*".

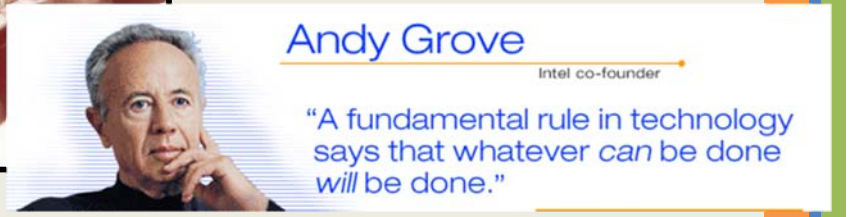

"Only the paranoid survive".

Moore received a B.S. degree in Chemistry from the University of California, Berkeley in 1950 and a **Ph.D. in Chemistry** and minor in Physics from Caltech in 1954.

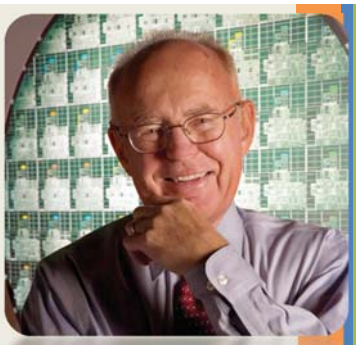

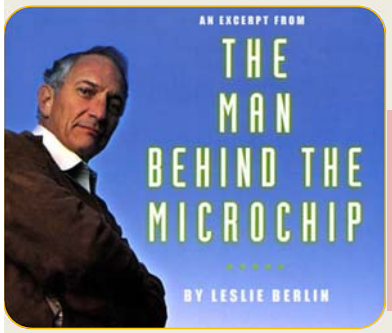

Noyce received his **Ph.D. in physics** from Massachusetts Institute of Technology in 1953. He is nicknamed "**the Mayor of Silicon Valley**"

Noyce was a mentor and father-figure to an entire generation of entrepreneurs, including Steve Jobs at Apple, Inc.

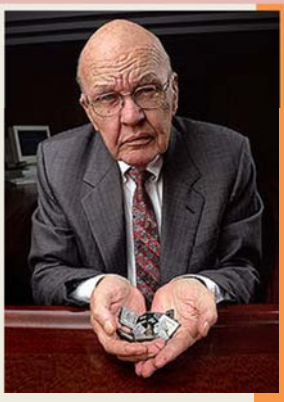

Noyce is credited (along with **Jack Kilby**) with the invention of the integrated circuit (IC).
(While Kilby's invention was **six months earlier**, neither man rejected the title of co-inventor.)

Jack Kilby was a Nobel Prize laureate in physics (2000):
Invention of the IC while working at Texas Instruments (TI).

Grove earned a Ph.D. in chemical engineering from the UC Berkeley.

He became Intel's president in 1979, its CEO in 1987, and its Chairman and CEO in 1997.

Grove is credited with **having transformed Intel from a manufacturer into one of the world's dominant producers of microprocessors**.

During his tenure as CEO, Grove oversaw a 4,500% increase in Intel's market capitalization making Intel the world's most valuable company.

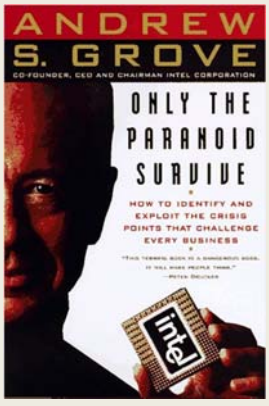

Grove was fiercely competitive, and he and the company became known for his guiding motto: "Only the paranoid survive".

Noyce was essentially anti-competitive: This difference in styles reputedly caused some degree of friction between Noyce and Grove.

Early Intel: 8008 and 8080

The **Intel 8008:** 8-bit byte-oriented microprocessor (April 1972).

The **Intel 8080** (April 1974):

- Running at 2 MHz (at up to 500,000 instructions per second), - Sometimes considered to be the first truly usable microprocessor.

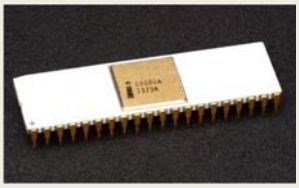

Shortly after the launch of the 8080, the Motorola 6800 competing design was introduced. Zilog introduced the Z80, which had a compatible machine-language instruction set and initially used the same assembly language as the 8080.

At Intel, the 8080 was followed by the compatible and electrically more elegant 8085, and later by the assembly language compatible 16-bit 8086 and then the 8/16-bit 8088, which was selected by IBM for its new PC to be launched in 1981.

The MITS Altair 8800 was a microcomputer design from 1975 based on the Intel 8080

Originally for a few hundred build-it-yourself kits to hobbyists (sold thousands in the first month)

Today the Altair is widely recognized as the spark that led to the microcomputer revolution of the next few years - First programming language for the machine was **Microsoft's founding product, Altair BASIC.**

A le ld be
inte by
Bill d no
BAS

The
mac

Alle SIC.

The shed,
but

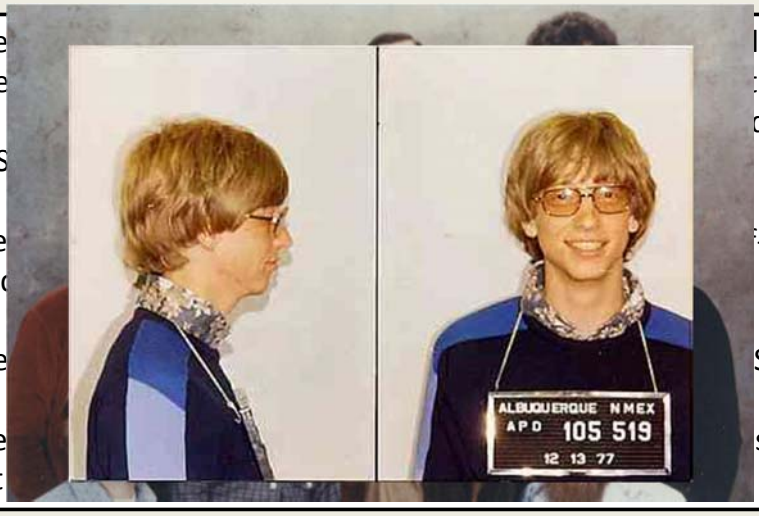

Apple was founded on **April 1, 1976** by Steve Jobs, Steve Wozniak, and Ronald Wayne to sell the Apple I personal computer kit.

They were hand-built by Steve Wozniak in the living room of Jobs' parents' home, and the Apple I was first shown to the public at the Homebrew Computer Club. Eventually 200 computers were built.

Steve Wozniak looked at the Intel 8080 chip (the heart of the Altair), but at **$179 decided he couldn't afford it**. Another chip, the Motorola 6800, interested Wozniak .

MOS Technology sold their 6502 chip (almost identical to the 6800) for $25, as opposed to the $175 Motorola 6800.

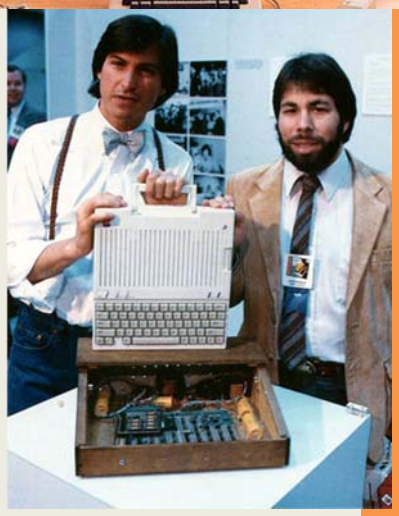

And… Now

The first Apple II computers went on sale on June 5, 1977] with a MOS Technology 6502 microprocessor running at 1 MHz, 4 KB of RAM

Throughout the 1980s and much of the 1990s, the Apple II was the *de facto standard computer* in American education


Welcome,
IBM.

Seric

Welcome to the most excitin
since the computer revolution be
And congratulations on you
Putting real computer powe
is already improving the way pe
communicate and spend their le
Computer literacy is fast bec
as reading or writing.
When we invented the first
we estimated that over 140,000,0
justify the purchase of one, if onl
Next year alone, we project
come to that understanding. Ove
of the personal computer will continue in logarithmic leaps.
We look forward to responsible competition in the massive
effort to distribute this American technology to the world.
And we appreciate the magnitude of your commitment.
Because what we are doing is increasing social capital
by enhancing individual productivity.
Welcome to the task.

apple

1981: Welcome IBM Seriously

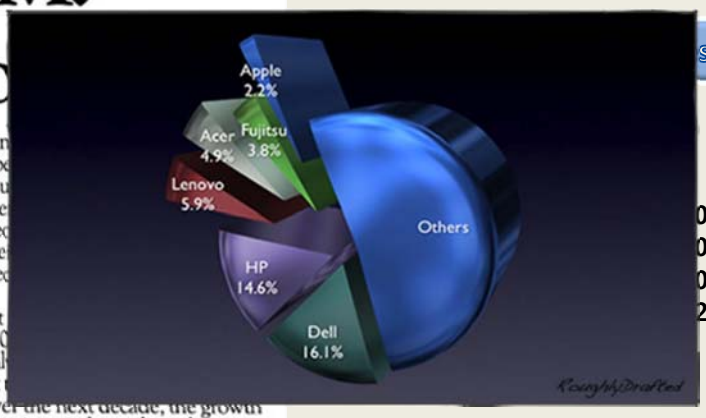

ssor + MS OS

	Introduced	CPU
0	Aug-81	8088
0	Mar-83	8088
0	Aug-84	80286
2	Sep-86	80286

Apple failed to make any progress in the 1990's

On July 3, 1991, IBM offered to help Apple finish Pink, its object oriented operating system for Jaguar, if Apple would adopt the PowerPC processor. Motorola was brought in to help manufacture the new processors, and the deal was sealed, creating the **Apple-IBM-Motorola (AIM)** alliance.

A 8MHz Intel 80286 Microprocessor

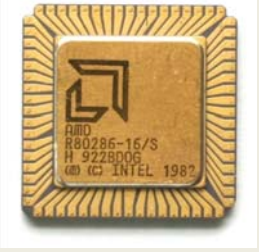

AMD 80286 (16 MHz version)

An 16-bit x86 microprocessor with 134,000 transistors (February 1, 1982)

It was employed for the IBM PC/AT, introduced in 1984, and then widely used in most PC/AT compatible computers until the early 1990s.

Intel 80386 DX rated at 16 MHz

The Intel 80386 (AKA the i386, or just 386)

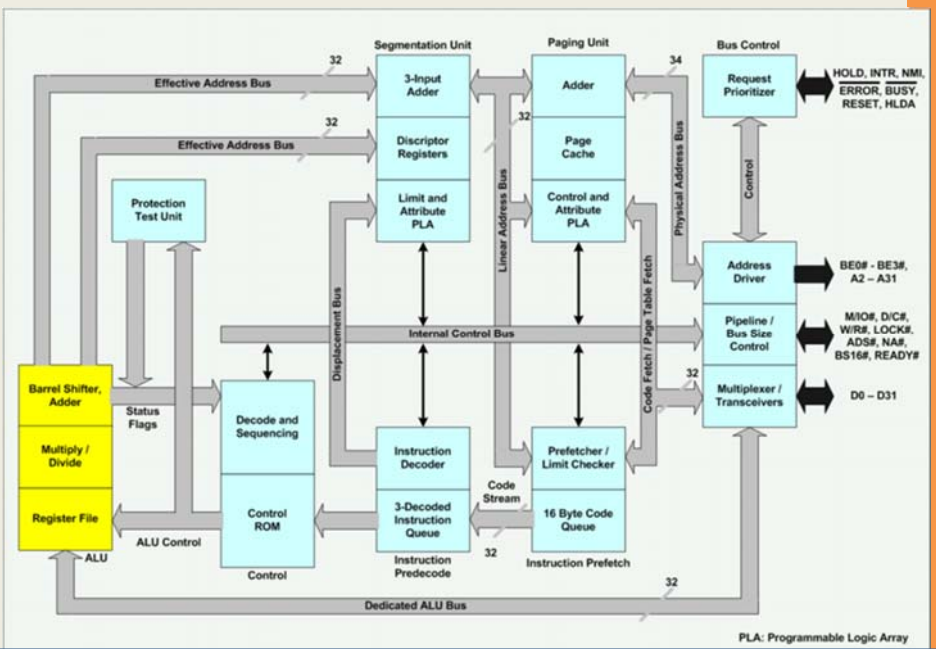

A 32-bit microprocessor (1985): the first versions had 275,000 transistors

The 32-bit extensions to the 8086 architecture, the 80386 instruction set

Still the denominator for all 32-bit x86 processors (x86, IA-32, or the i386-architecture)

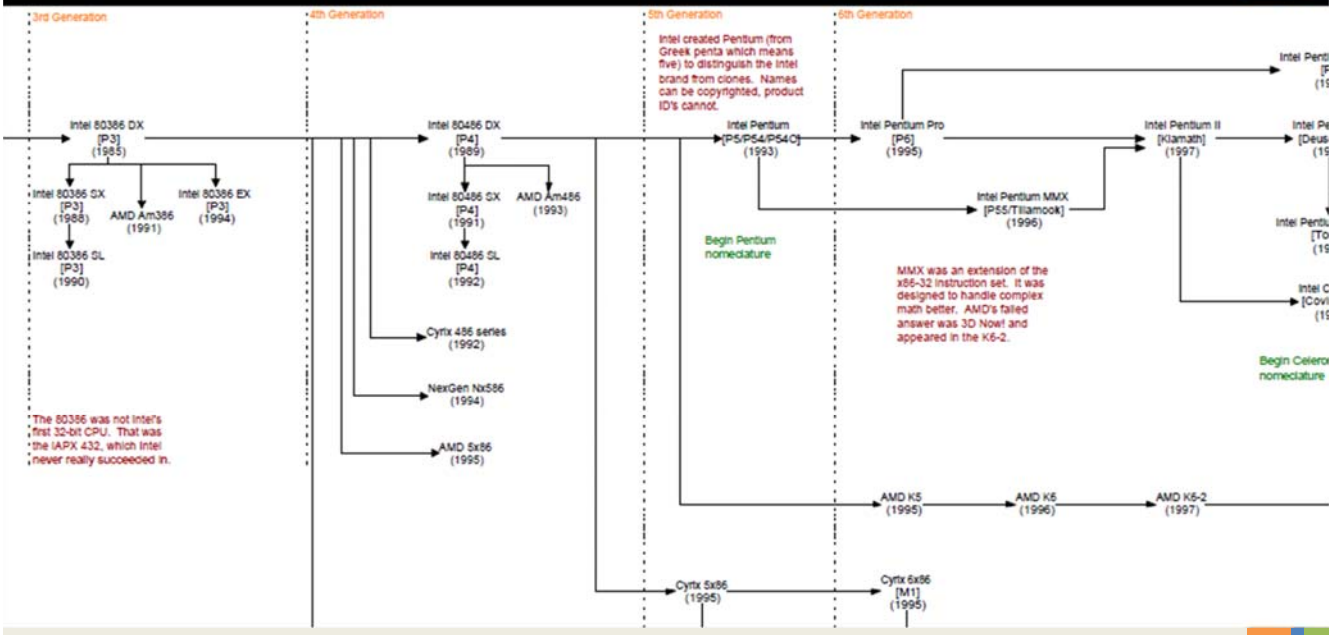
3rd Generation
4th Generation
5th Generation
6th Generation
Intel created Pentium (from Greek penta which means five) to distinguish the Intel brand from clones. Names can be copyrighted, product ID's cannot.
Intel 80386 DX [P3] (1985)
Intel 80386 SX [P3] (1988)
AMD Am386 (1991)
Intel 80386 EX [P3] (1994)
Intel 80386 SL [P3] (1990)
Intel 80486 DX [P4] (1989)
Intel 80486 SX [P4] (1991)
AMD Am486 (1993)
Intel 80486 SL [P4] (1992)
Cyrix 486 series (1992)
NexGen Nx586 (1994)
AMD 5x86 (1995)
Intel Pentium [P5/P54/P54C] (1993)
Intel Pentium Pro [P6] (1995)
Intel Pentium MMX [P55/Tillamook] (1996)
Intel Pentium II [Klamath] (1997)
Begin Pentium nomeclature
MMX was an extension of the x86-32 instruction set. It was designed to handle complex math better. AMD's failed answer was 3D Now! and appeared in the K6-2.
The 80386 was not Intel's first 32-bit CPU. That was the iAPX 432, which Intel never really succeeded in.
AMD K5 (1995)
AMD K6 (1996)
AMD K6-2 (1997)
Cyrix 5x86 (1995)
Cyrix 6x86 [M1] (1995)

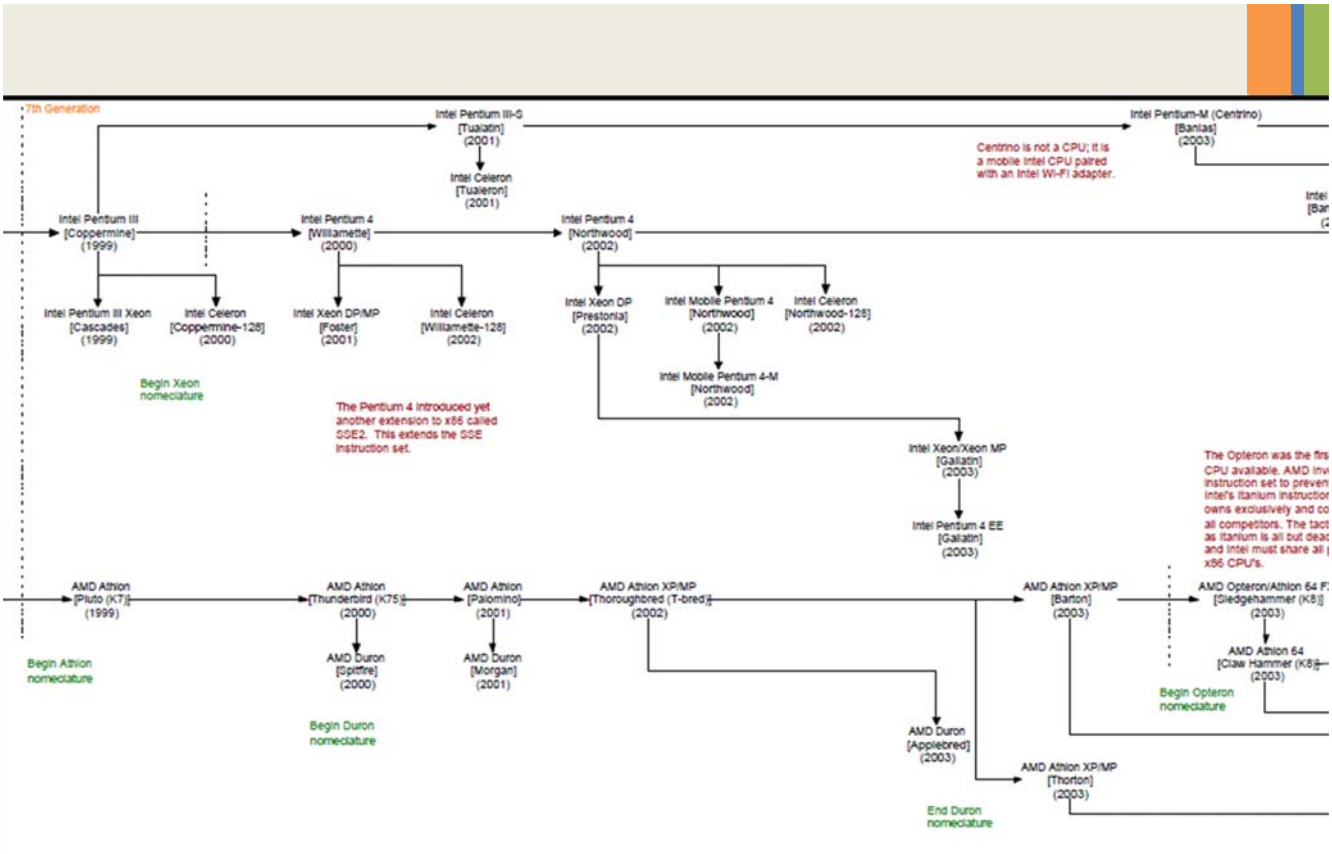
7th Generation
Intel Pentium III-S [Tualatin] (2001)
Intel Celeron [Tualeron] (2001)
Intel Pentium-M (Centrino) [Banias] (2003)
Centrino is not a CPU; it is a mobile Intel CPU paired with an Intel Wi-Fi adapter.
Intel Pentium III [Coppermine] (1999)
Intel Pentium 4 [Willamette] (2000)
Intel Pentium 4 [Northwood] (2002)
Intel Pentium III Xeon [Cascades] (1999)
Intel Celeron [Coppermine-128] (2000)
Intel Xeon DP/MP [Foster] (2001)
Intel Celeron [Willamette-128] (2002)
Intel Xeon DP [Prestonia] (2002)
Intel Mobile Pentium 4 [Northwood] (2002)
Intel Celeron [Northwood-128] (2002)
Intel Mobile Pentium 4-M [Northwood] (2002)
Begin Xeon nomeclature
The Pentium 4 introduced yet another extension to x86 called SSE2. This extends the SSE instruction set.
Intel Xeon/Xeon MP [Gallatin] (2003)
Intel Pentium 4 EE [Gallatin] (2003)
AMD Athlon [Pluto (K7)] (1999)
AMD Athlon [Thunderbird (K75)] (2000)
AMD Athlon [Palomino] (2001)
AMD Athlon XP/MP [Thoroughbred (T-bred)] (2002)
AMD Athlon XP/MP [Barton] (2003)
AMD Opteron/Athlon 64 FX [Sledgehammer (K8)] (2003)
AMD Athlon 64 [Claw Hammer (K8)] (2003)
Begin Athlon nomeclature
AMD Duron [Spitfire] (2000)
AMD Duron [Morgan] (2001)
Begin Duron nomeclature
Begin Opteron nomeclature
AMD Duron [Applebred] (2003)
AMD Athlon XP/MP [Thorton] (2003)
End Duron nomeclature

Pentium 4 ….

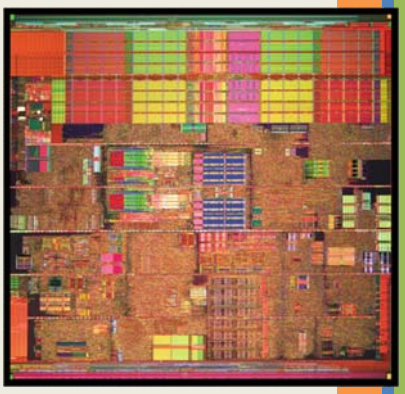

Intel's single core microprocessor:
The first model: 1.5GHz, November, 2000.

It includes several important new features and innovations that will allow the Intel Pentium 4 processor to deliver industry-leading performance for the next several years.

Produced: 2000 ~ 2008
Max. CPU rate: 1.3 GHz to 3.8 GHz

At the launch of the Pentium 4, Intel stated NetBurst-based processors were expected to scale to **10 GHz** (which should be achieved over several fabrication process generations). However, the NetBurst microarchitecture ultimately hit a frequency ceiling far below that expectation – the fastest clocked NetBurst-based models reached a peak clock speed of 3.8 GHz.

Intel had not anticipated a rapid upward scaling of **transistor power leakage** that began to occur as the chip reached the 90 nm process node and smaller. This new power leakage phenomenon, along with the standard thermal output, **created cooling and clock scaling problems** as clock speeds increased.

The Pentium III was eventually superseded by the Pentium 4, but its **Tualatin** core also served as the **basis for the Pentium M CPUs.**

Subsequently, it was the P-M microarchitecture of Pentium M branded CPUs (not the NetBurst found in Pentium 4 processors), **that formed the basis for Intel's energy-efficient Intel Core microarchitecture of CPUs branded Core 2 and Xeon**.

August 8, 2008: The End of Pentium 4

Santa Clara (CA) – Intel today officially announced the Xeon X5365 – a quad-core processor that so far only has been available in limited quantities. The company also quietly announced that it has begun phasing out all remaining Pentium 4 and Pentium D processors.

Intel's 9th Generation Microarchitecture

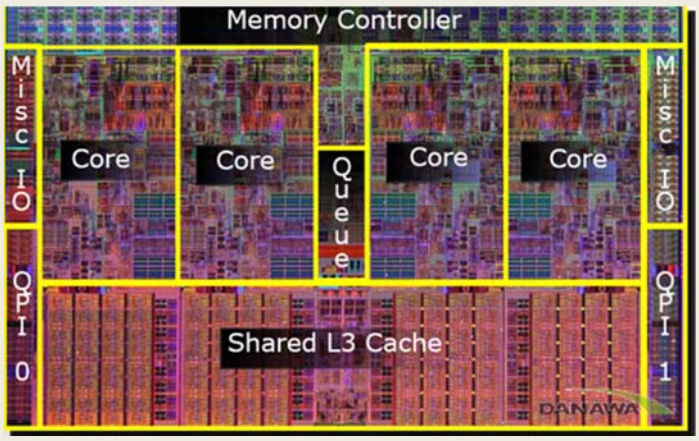

GPU Architecture: NVIDIA Tesla

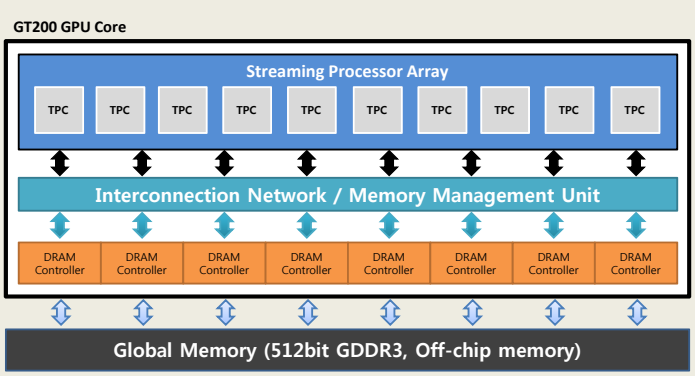

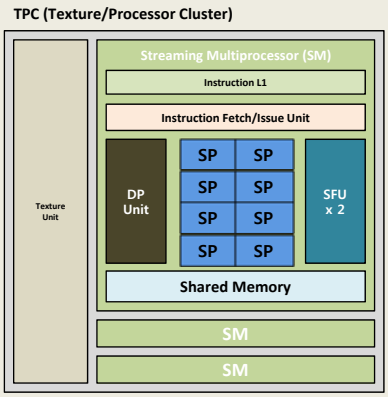

- Streaming Multiprocessor(SM)
  - Including eight scalar processors
  - Instruction issuing, shared memory and cache control
- Scalar Processor (SP)
  - up to 240 cores, integer computation unit

GPU Architecture: NVIDIA Fermi

- Configurable L1 and Unified L2 cache
  - a true cache hierarchy for load/store operations
- Unified CUDA core
  - a fully pipelined ALU and FPU
- Up to 512 CUDA cores
- Improve double precision floating point
  - 4x increase over Tesla architecture
- 1.5 TFLOPs computing power

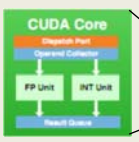

CPU vs. GPU - Vector Addition

CPU

```
for(int i=0; i<N ; i++){
    c[i] = a[i] + b[i];
}
```

GPU

```
c[threadIdx.x] = a[threadIdx.x]
+ b[threadIdx.x];
```

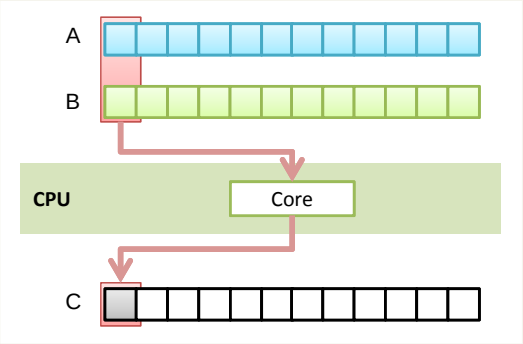

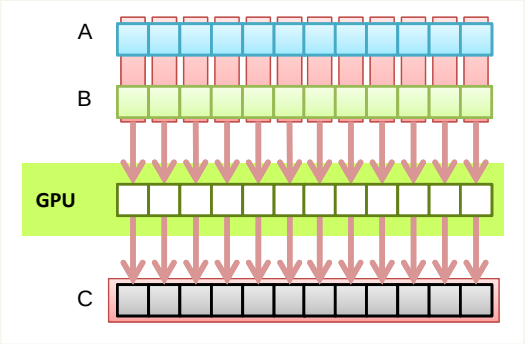

GPGPU Programming Model: OpenCL

- Open Computing Language (OpenCL)
- Open standard framework for parallel programming of heterogeneous systems
  - Developed by Apple Inc. and refined by technical teams at AMD, IBM, Intel, and NVidia.
  - Industry standard for parallel programming across CPUs and GPUs
  - Implementations are all based on the LLVM (Low Level Virtual Machine) Compiler technology and use the Clang Compiler as its frontend
  - OpenCL 1.1 was released on June 14, 2010
  - Diverse industry participation

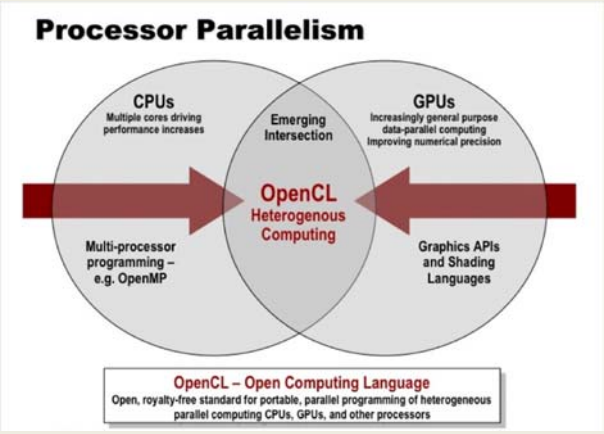

Heterogeneous Chip Multiprocessor

- General purpose processor which integrates multiple processing units on the same die
  - High programmability and performance
  - Prioritizes bandwidth over latency
  - High performance core has a fast local-access memory or large internal registers
  - High performance core favors SIMD computations

- Representative Processors
  - IBM Cell BE (Cell Broadband Engine)
  - Intel Sandy Bridge
  - AMD Fusion APU

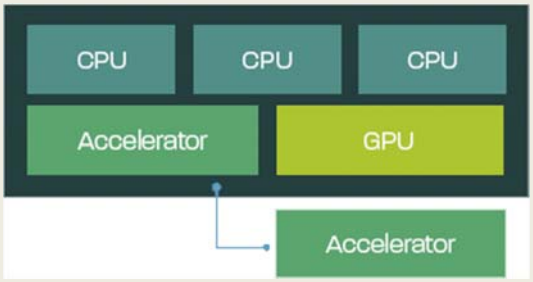

Cell BE Architecture

- Consisted of one Power Processor Element (PPE), and eight co-processors called Synergistic Processing Elements (SPEs)
- PPE
  - 64-bit PowerPC Architecture
  - Handles most of computational workload
- SPEs
  - A RISC processor with 128-bit SIMD organization
  - 256 KB SRAM Local Storage (LS) and no hardware cache
  - High performance data transfer asynchronously and in parallel with data processing on the PPE and SPEs
- EIB (Element Interconnect Bus)
  - A circular ring bus comprising of four 16B-wide unidirectional channels which counter-rotate in pairs
  - Connecting the PPE, input/output elements, and the SPEs

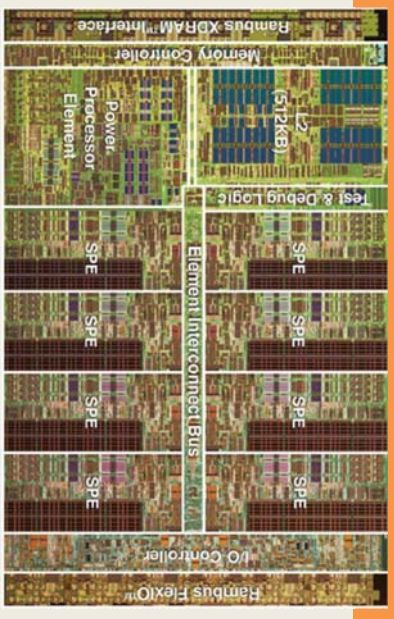

Intel Sandy Bridge Architecture

- Cores and Graphics are Integrated on a Single Chip
  - First released on January 2011.01
- AVX (Advanced Vector eXtension)
  - Featuring a 256-bit instruction set with wider vectors, new extensible syntax, and rich functionality
- Ring Bus Interconnect
  - Fast access by cores and graphics to shared data in the last-level cache accelerates graphics processing
  - Shared last-level cache remains “sliced” (NUMA), with slices distributed among cores
- op Cache (CISC to RISC instructions)
  - Increasing performance by more consistently delivering uops to the back-end and eliminating various bubbles in the fetch and decode process
- OpenCL Support
  - Intel has decided to support OpenCL (Sep. 13, 2010)

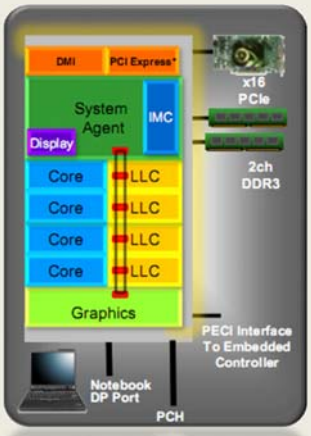

Smartphone

by IBM in 1992

Microsoft Windows CE

in 2001

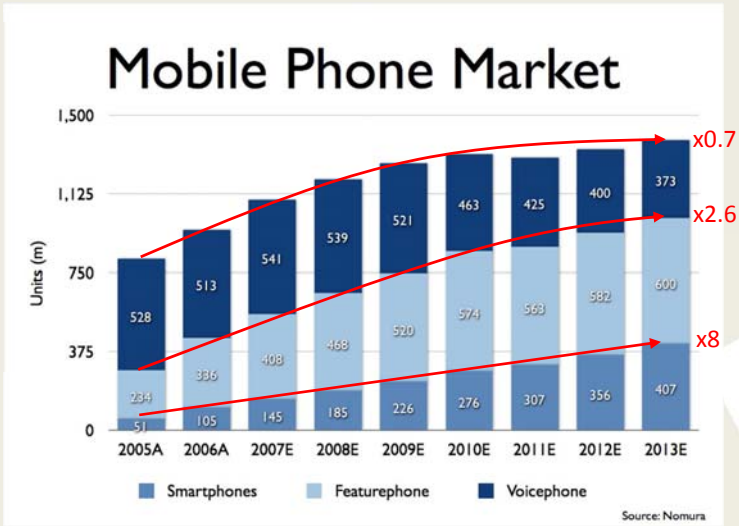

in 2010

Blackberry in 2002

iPhone in 2007

Android in 2010

Trends: High-Performance Mobile CPU + GPU

The user's demands for mobile systems have been increased and various mobile applications have been developed.

Mobile GPUs

Mali- 55
1 fragment processor
1M tri/s, 100Mpix/s

PowerVR MBX
3.4-7.4M tri/s, 270-600Mpix/s

Innovation

Latest Mobile GPUs

Adreno GPU (AMD, Qualcomm)
Hardware-accelerated SVG and Adobe Flash performance

PowerVR SGX (Imagination Tech.)
Multi-threaded shader engine
7-140M tri/s, 250M-4Gpix/s

GeForce GPU 400 (Tegra 2)
Cortex-A9 MP core, OpenGL ES 2.0
Programmable pixel shader

Influence

Desktop GPUs

Tesla S1070
960 cores , 4.1T(single)-345G(double)FLOPS

FireStream9270
800 cores, 1.2T(single)-240G(double)FLOPS

Future trends for high performance mobile applications:
Parallel computing on heterogeneous MP-SoC with GPGPU

Multi-Core on Mobile Devices

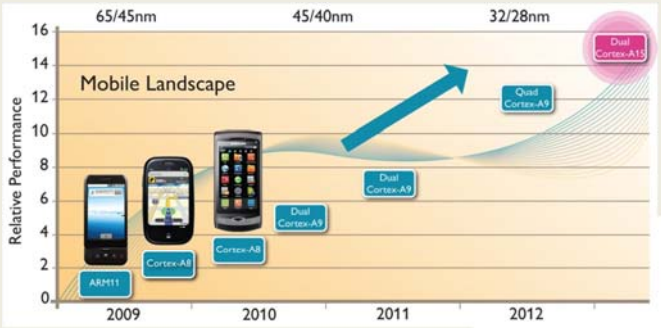

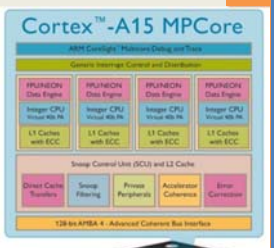

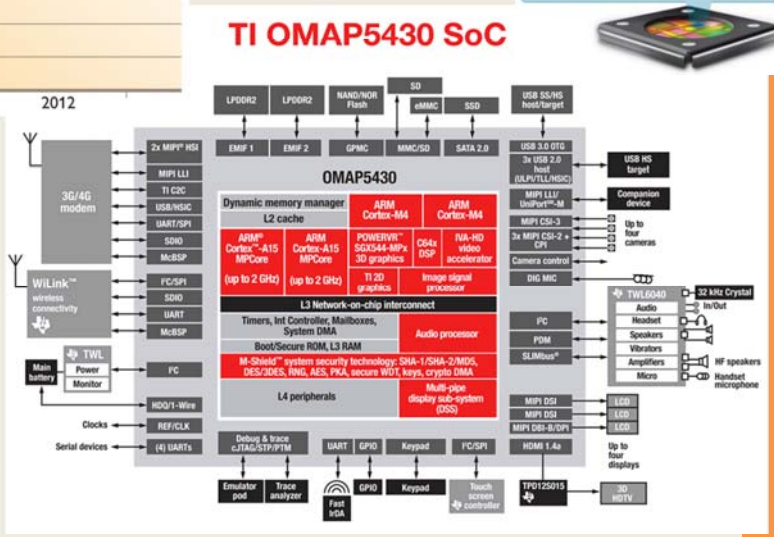

- **Further enhancements to simplify and broaden the adoption of multicore solutions**
- **Mobile configurations of the Cortex-A15 processor will deliver over five times the performance of today's advanced smartphones**

Tegra

- NVIDIA Tegra 2

- CPUs and GPU on the same die
- Ultra-low power (ULP) GeForce GPU
- Dual-core ARM Cortex-A9 processor
- The first mobile CPU with out of order execution

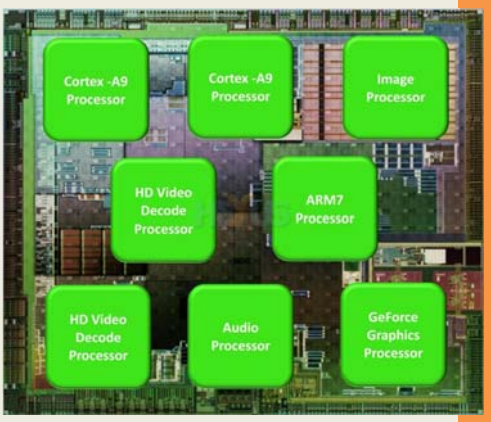

- ARM 7 processor
- Supports battery management functions
- Ultra-low power (ULP) GeForce GPU

Snapdragon

- Qualcomm Snapdragon (3rd Generation)

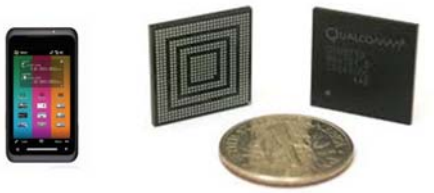

- Scorpion dual-core CPU, up to 1.5GHz
- Full support for websites based on Flash 10 and WebGL
- Integrated 3G broadband connectivity
- 45 nm semiconductor technology → Low-power

- Advanced mobile graphics
- 1080p video recording and playback, up to 30fps
- OpenGL ES2.0, OpenVG 1.1 acceleration
- Built-in 8th-Gereration GPS engine
- Support Wi-Fi and Bluetooth

Hummingbird & Orion

- **Samsung Hummingbird**

  - Used in Galaxy S, Galaxy Tab, and Galaxy Player
  - 45 nm ARM Cortex-A8 processor @~1.2 GHz
  - Powerful 3D rendering, video decoding, and image processing with Power VR SGX540 GPU

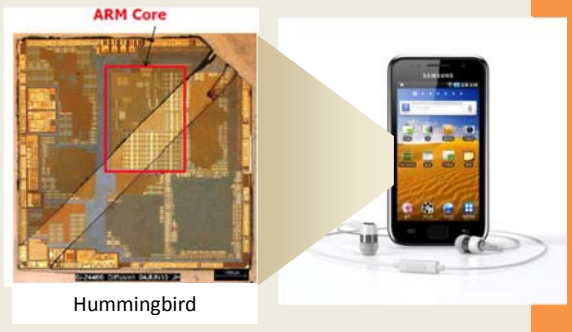

Hummingbird

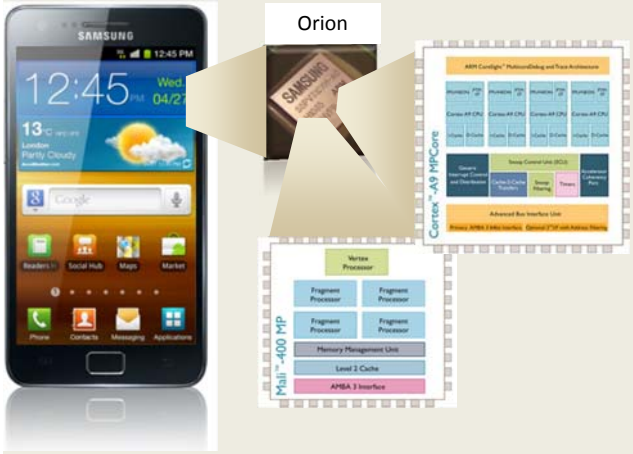

- **Samsung Orion (Exynos)**

  - Used in Galaxy S2
  - 45 nm dual-core ARM Cortex-A9 processor
  - Clock frequency up to 1.2 GHz
  - Low-power design, DDR3 support, and audio DSP

A4

- Designed by Apple and manufactured by Samsung
- ARM Cortex-A8 CPU with a Power VR GPU
- Used in iPhone 4, iPad tablet, iPod Touch, and Apple TV
- PowerVR SGX 535 GPU
- Samsung's 45 nm silicon chip fabrication process
- Low-power design

A5

Display	• 9.7-inch (diagonal) LED-backlit glossy widescreen Multi-Touch display with IPS technology • 1024-by-768-pixel resolution at 132 pixels per inch (ppi) • Fingerprint-resistant oleophobic coating • Support for display of multiple languages and characters simultaneously
Chip	1GHz dual-core Apple A5 custom-designed, high-performance, low-power system-on-a-chip

Atom & Nano X2

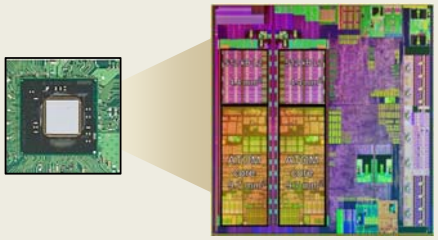

- **Intel Atom (Pineview)**
  - **Dual-core Atom processor released on Q3' 2010**
  - **Integrated the graphics, display, and memory controller**
  - **Advanced Smart Cache**

- **VIA Nano X2**
  - **Dual-core Isaiah processor (x86-64 architecture)**
  - **Released on January, 2011 to rival Intel Atom**
  - **40nm fabrication process**
  - **Advanced 64-bit, superscalar architecture**

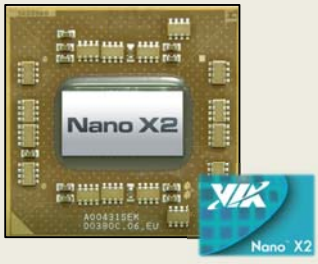

✓ ***Diverse Netbook Operating Systems***

- Open projects for supporting these power-efficient processor
  - Android x86-port, Chromium OS, MeeGo, Ubuntu Netbook edition, and so on
- Mobile devices in emerging market and multi-core support

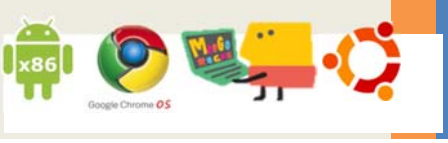

Low Power Data Center
ARM vs. Atom (x86)

ARM

"Project Denver extends the range of ARM system upward to PCs, data center servers, and supercomputers."

Nvidia Developing ARM Processor for Servers

January 6th, 2011 : Rich Miller

- NVIDIA Blog – Chief Scientist Bill Dally: "ARM is already the standard architecture for mobile devices. Project Denver extends the range of ARM systems upward to PCs, data center servers, and supercomputers. ARM's modern architecture, open business model, and vibrant eco-system have led to its pervasiveness in cell phones, tablets, and other embedded devices. Denver is the catalyst that will enable these same factors to propel ARM to become pervasive in higher-end systems. Denver frees PCs, workstations and servers from the hegemony and inefficiency of the x86 architecture."

Intel SCC & MIC

- **SCC (Single-chip Cloud Computer)**
  - 48 cores on a silicon CPU chip
  - 24 tiles with two IA cores per tile
  - 24-router mesh network
  - 4 integrated DDR3 memory controller
  - Each core can run a separate OS and software stack over a packet-based network

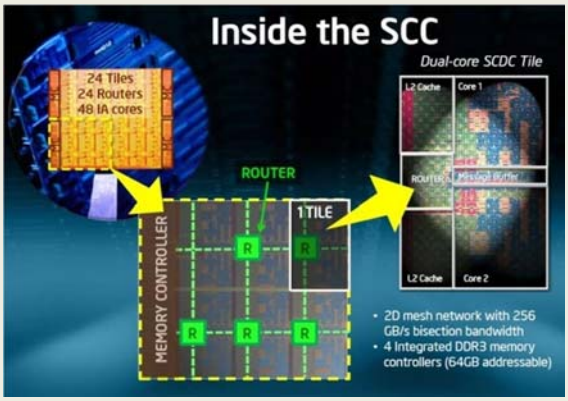

- **MIC (Many Integrated Core) Architecture**
  - Symmetric multithreaded CPU
  - Many-core array chip on a single die, integrating 10~100 cores
  - Knights Corner (the first MIC product) 22 nm manufacturing process and 50 Intel processing cores on a single chip