
IIT BOMBAY
CARNEGIE MELLON UNIVERSITY
PITTSBURGH PENNSYLVANIA 1900
From Mainframe to Smartphone: What an Amazing Trip It's Been!
TEXAS INSTRUMENTS
digital
digital
Dileep Bhandarkar, Ph. D.
IEEE Fellow
Computer History Museum
21 August 2014
intel
intel
Microsoft
QUALCOMM

Disclaimer

The opinions expressed here are my own and
may be a result of the way in which
my highly disorganized and somewhat forgetful mind interprets
a particular situation or concept.

They are not approved or authorized by my current or past employers, family, or friends.

If any or all of the information or opinions found here does accidentally
offend, humiliate or hurt someone's feelings,
it is entirely unintentional.

“Come on the amazing journey
And learn all you should know.” – The Who

The Stops Along My Journey

- 1970: B. Tech, Electrical Engineering (Distinguished Alumnus)
  - Indian Institute of Technology, Bombay
- 1973: PhD in Electrical Engineering
  - Carnegie Mellon University
  - Thesis: Performance Evaluation of Multiprocessor Computer Systems
- 4 years - Texas Instruments
  - Research on magnetic bubble & CCD memory, Fault Tolerant DRAM
- 17.5 years - Digital Equipment Corporation
  - Processor Architecture and Performance
- 12 years - Intel
  - Performance, Architecture, Strategic Planning
- 5.5 years - Microsoft
  - Distinguished Engineer, Data Center Hardware Engineering
- Since January 2013 – Qualcomm Technologies Inc
  - VP Technology

"Follow the path of the unsafe, independent thinker. Expose your ideas to the danger of controversy. Speak your mind and fear less the label of ''crackpot'' than the stigma of conformity." – Thomas J. Watson

Directions | Contact Us

COMPUTER HISTORY MUSEUM

The Silicon Engine:
A Timeline of Semiconductors in Computers

MOORE'S LAW "Transistor density on integrated circuits doubles about every two years." *

1950s	1960s	1970s	1980s	1990s	2000s
Silicon Transistor	TTL Quad Gate	8-bit Microprocessor	32-bit Microprocessor	32-bit Microprocessor	64-bit Microprocessor
			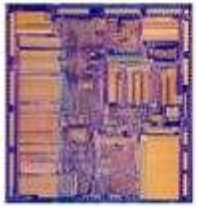		
1 Transistor	**16 Transistors**	**4500 Transistors**	**275,000 Transistors**	**3,100,000 Transistors**	**592,000,000 Transistors**

Microelectronic silicon computer "chips" have grown in capability from a single transistor in the 1950s to hundreds of millions of transistors per chip on today's microprocessor and memory devices. From the first documented semiconductor effect in 1833 to the transition from transistors to integrated circuits in the 1960s and 70s, this website explores key milestones in the development of these extraordinary engines that power the computing and communications revolution of the information age.

*Source: "Moore's Law: Raising the Bar" (Intel Corporation 2005)

Photo credits: Fairchild Camera and Instrument Corporation, Intel Corporation (Note that images are not to scale)

1958: Jack Kilby's Integrated Circuit

SSI -> MSI -> LSI -> VLSI -> OMGWLSI

What is Moore's Law?

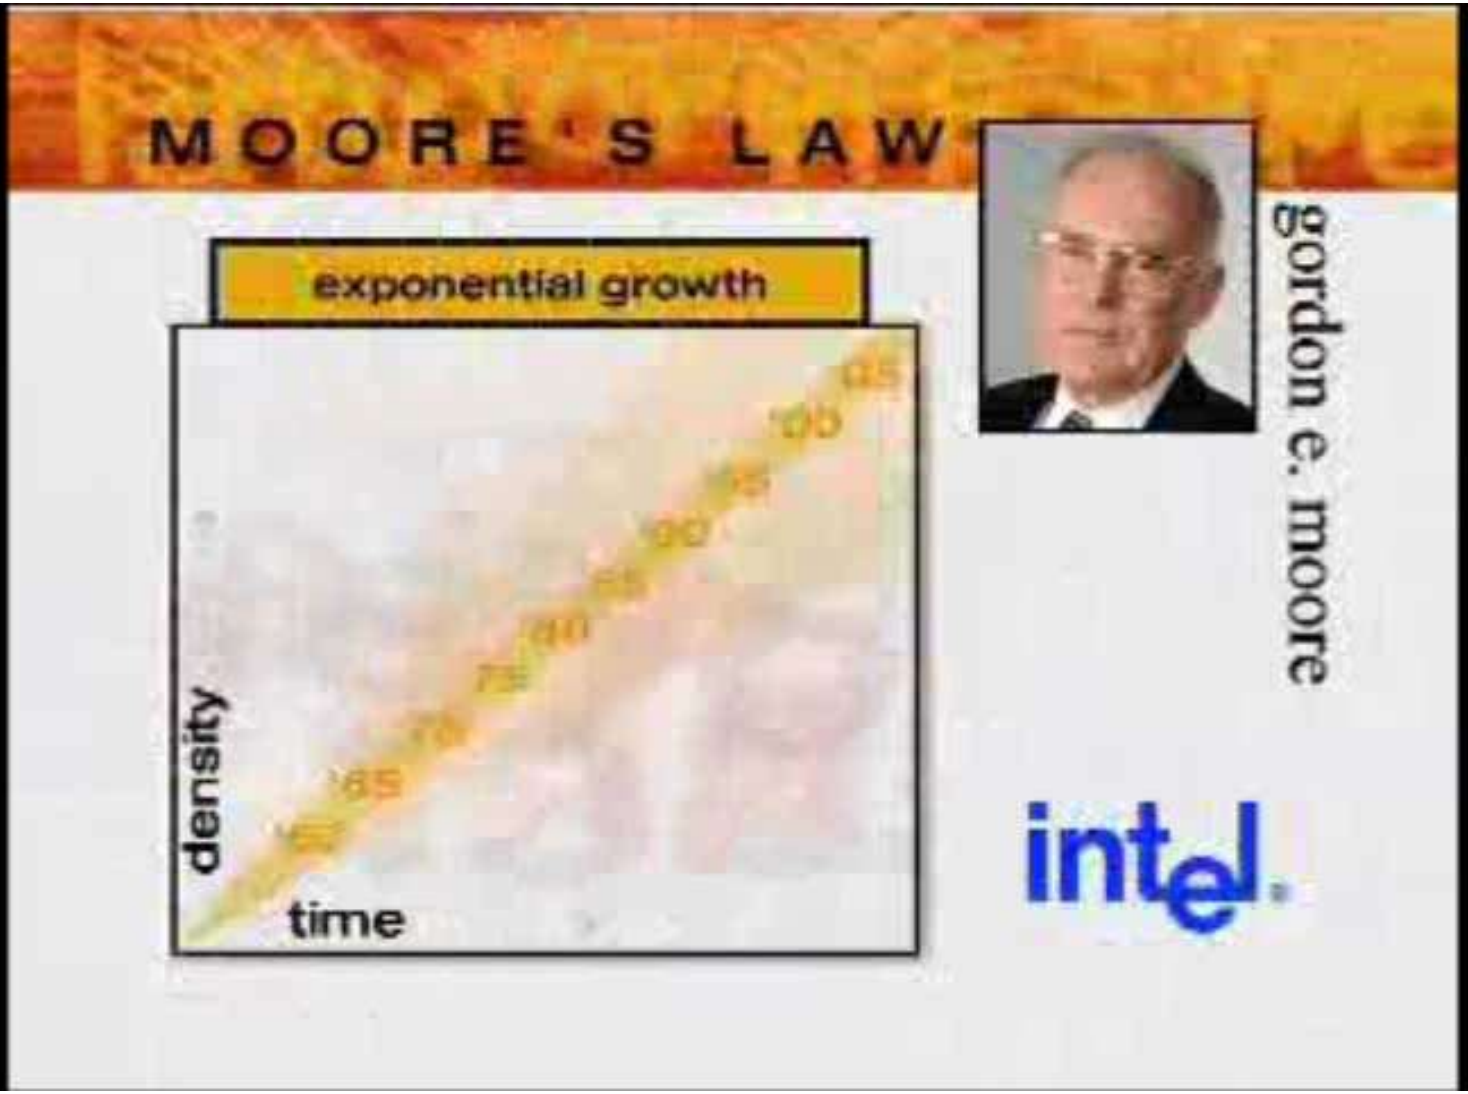


1970 – 73: Graduate School
Carnegie Mellon University
Pittsburgh, PA

1971: 4004 Microprocessor

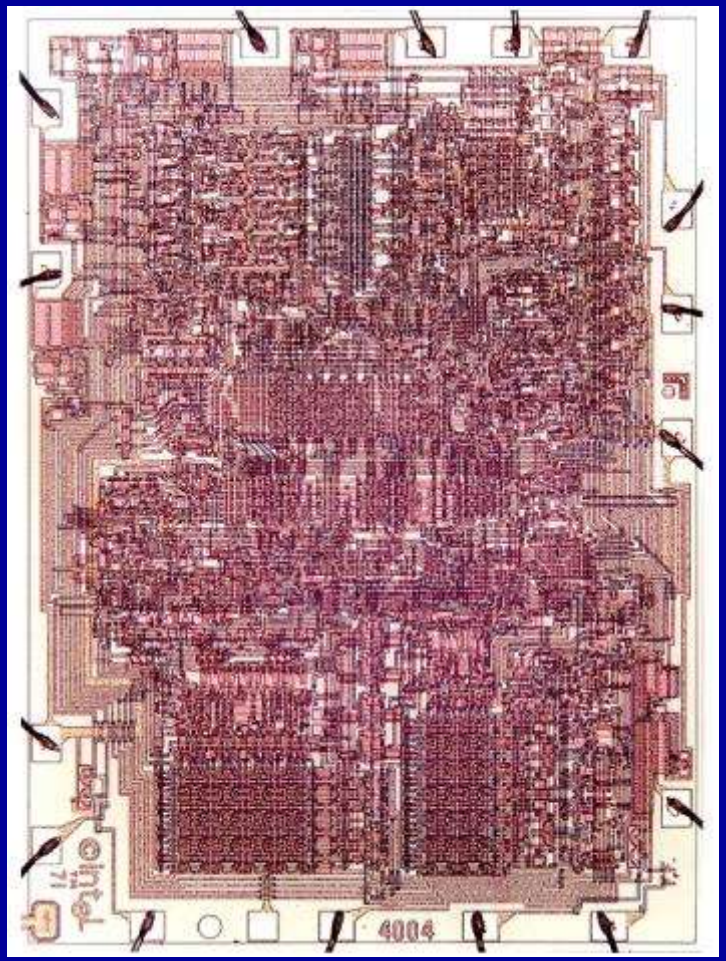

- The 4004 was Intel's first microprocessor. This breakthrough invention powered the Busicom calculator and paved the way for embedding intelligence in inanimate objects as well as the personal computer.

Introduced November 15, 1971
108 KHz, 50 KIPs , 2300 10μ transistors

1971: 1K DRAM

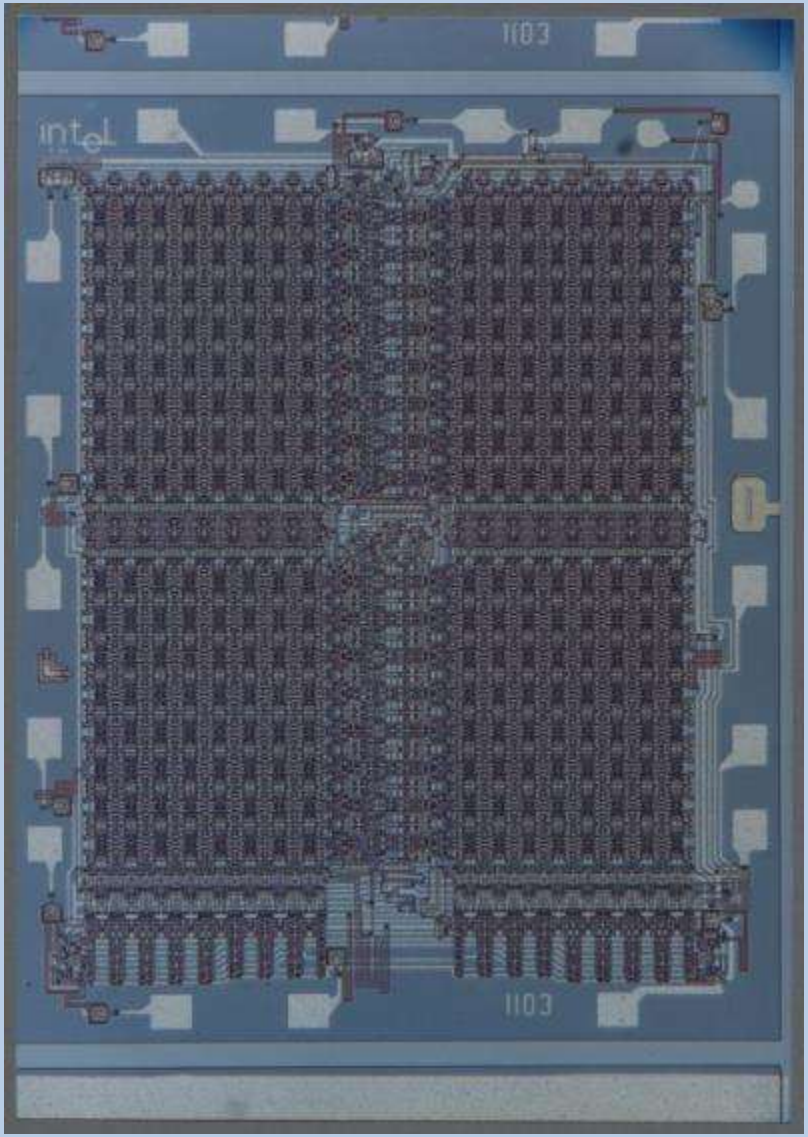

Intel® 1103 DRAM Memory

- Intel delivered the 1103 to 14 of the 18 leading computer manufacturers.
- Since the production costs of the 1103 were much lower than the costs of a magnetic core memory the market developed rapidly, becoming the world's best selling memory chip and was finally responsible for the obsolescence of magnetic core memory.

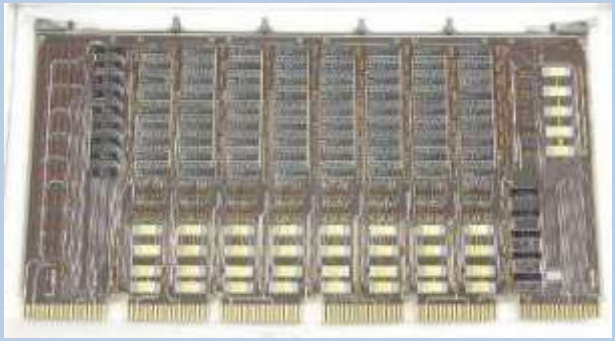

DRAM Memory Board

Core Memory

IBM 360/67 and Univac 1108 at CMU in 1970

- The S/360-67 operated with a basic internal cycle time of 200 nanoseconds and a basic 750 nanosecond magnetic core storage cycle
- Dynamic Address Translation (DAT) with support for 24 or 32-bit virtual addresses using segment and page tables (up to 16 segments each containing up to 256 x 4096 byte pages)

Snow White (IBM) and the Seven Dwarfs (RCA, GE, Burroughs, Univac, NCR, Control Data, Honeywell)

DEC PDP-10

Sept 1973: Mission Accomplished

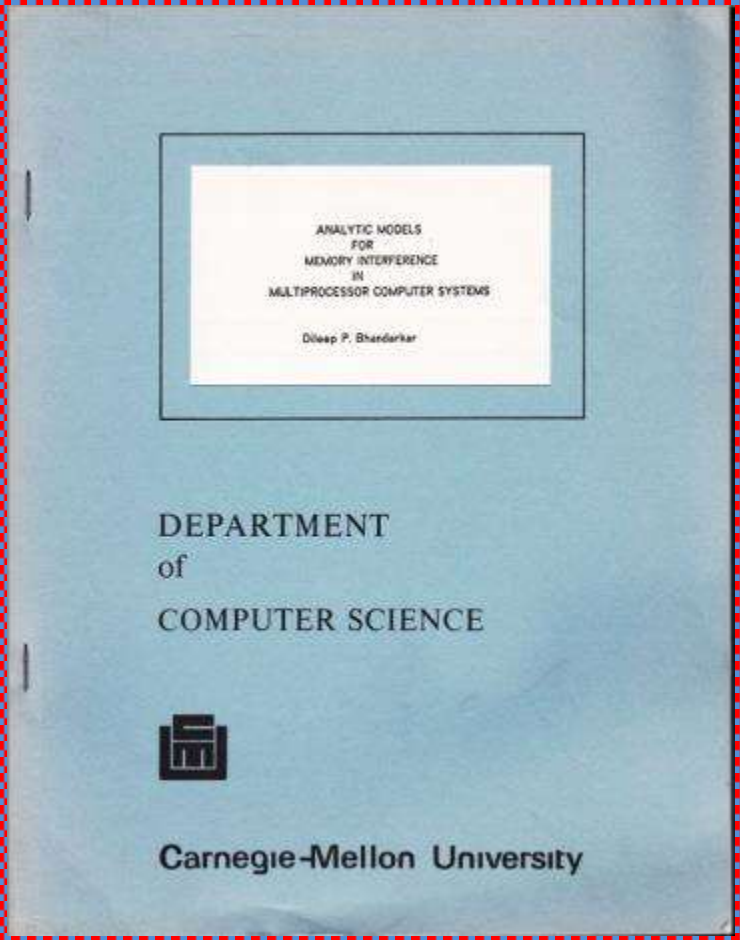

ANALYTIC MODELS
FOR
MEMORY INTERFERENCE
IN
MULTIPROCESSOR COMPUTER SYSTEMS

Dileep P. Bhandarkar

DEPARTMENT
of
COMPUTER SCIENCE

Carnegie-Mellon University

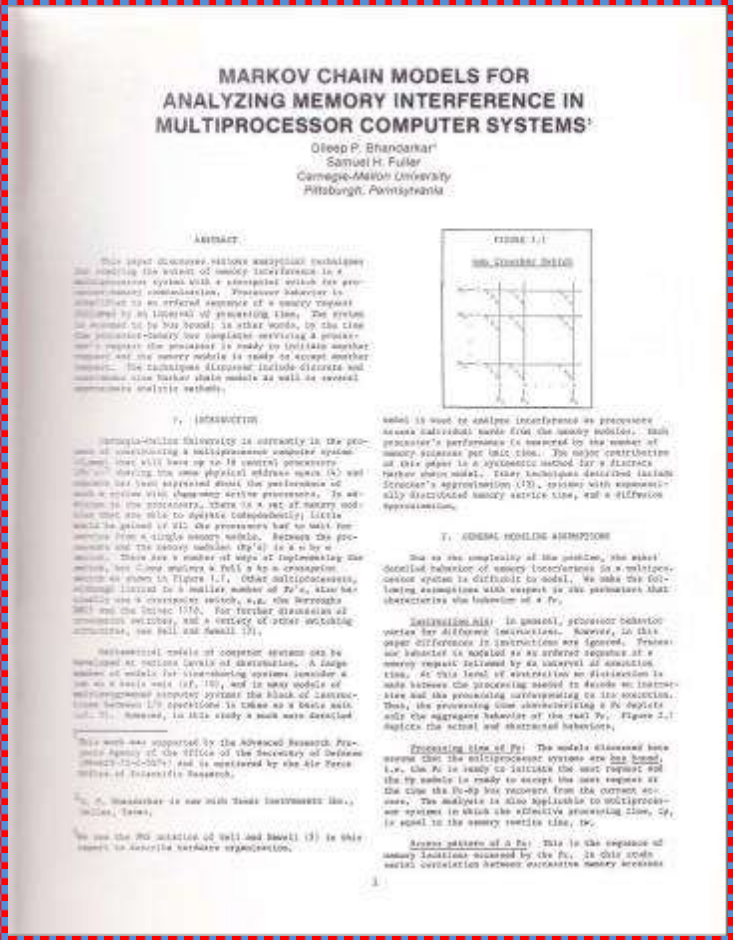

MARKOV CHAIN MODELS FOR
ANALYZING MEMORY INTERFERENCE IN
MULTIPROCESSOR COMPUTER SYSTEMS*

Dileep P. Bhandarkar
Samuel H. Fuller
Carnegie-Mellon University
Pittsburgh, Pennsylvania

Created with RUNOFF (XOFF) and printed on Xerox Graphics Printer prototype connected to a DEC PDP-11/20 running printer software developed by Chuck Geschke.

1st paper presented at the First Annual Symposium on Computer Architecture (later named ISCA) in Dec 1973 after Opening Keynote by Maurice Wilkes!

Oct 1973
First Job in Dallas, Texas

1973: 4K DRAM

22 pin package

TI and Intel used 22 pins for their competing, next-generation 4K devices in 1973. But Mostek soon dominated the 4K market by squeezing it into a 16 pin package using an address multiplexing scheme, which was a revolutionary approach that reduced cost and board space. By 1976 everyone adopted Mostek's approach for 16K and larger DRAMs.

16 pin package with RAS/CAS

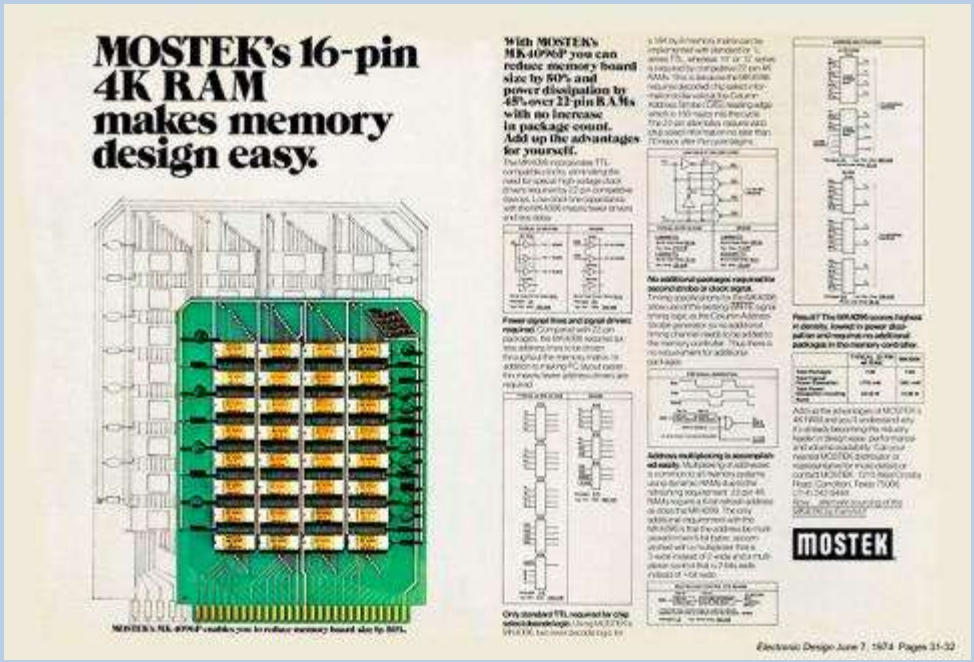

Texas Instruments Adventure

- Magnetic Bubble Memory
- Fault Tolerant DRAM

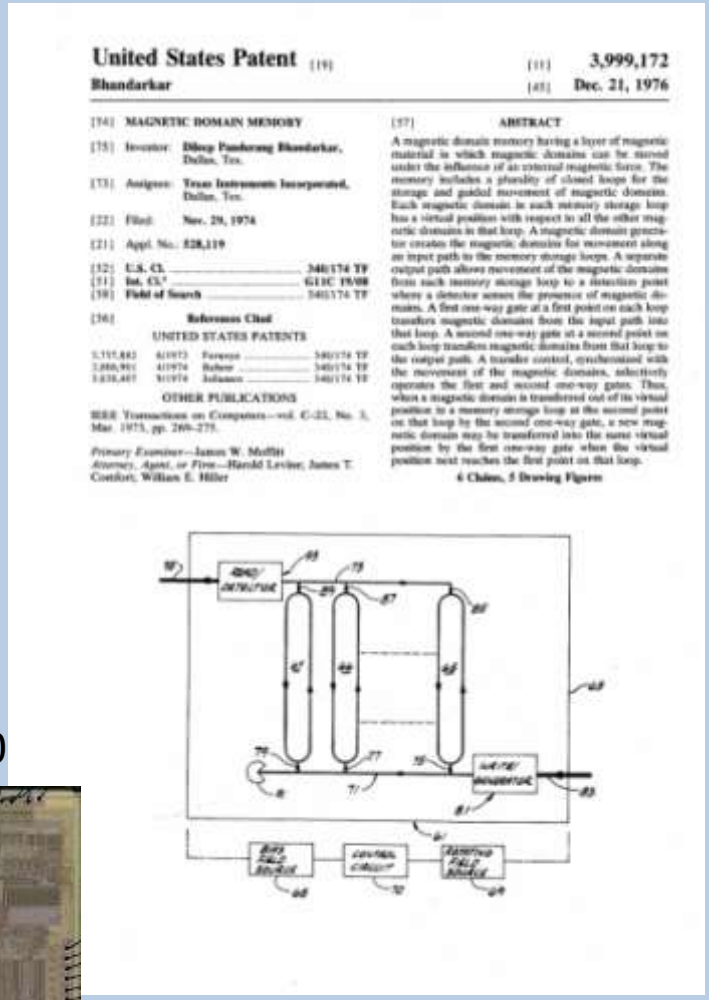

United States Patent [19]
Bhandarkar
[11] 3,999,172
[45] Dec. 21, 1976

[54] MAGNETIC DOMAIN MEMORY

[75] Inventor: Dileep Pandurang Bhandarkar, Dallas, Tex.

[73] Assignee: Texas Instruments Incorporated, Dallas, Tex.

[22] Filed: Nov. 29, 1974

[21] Appl. No.: 528,119

[52] U.S. Cl. ..... 340/174 TF
[51] Int. Cl.² G11C 19/08
[58] Field of Search ..... 340/174 TF

[56] References Cited
UNITED STATES PATENTS
OTHER PUBLICATIONS

Primary Examiner—James W. Moffitt
Attorney, Agent, or Firm—Harold Levine; James T. Comfort; William E. Hiller

[57] ABSTRACT

A magnetic domain memory having a layer of magnetic material in which magnetic domains can be moved under the influence of an external magnetic force. The memory includes a plurality of closed loops for the storage and guided movement of magnetic domains. Each magnetic domain in each memory storage loop has a virtual position with respect to all the other magnetic domains in that loop. A magnetic domain generator creates the magnetic domains for movement along an input path to the memory storage loops. A separate output path allows movement of the magnetic domains from each memory storage loop to a detection point where a detector senses the presence of magnetic domains. A first one-way gate at a first point on each loop transfers magnetic domains from the input path into that loop. A second one-way gate at a second point on each loop transfers magnetic domains from that loop to the output path. A transfer control, synchronized with the movement of the magnetic domains, selectively operates the first and second one-way gates. Thus, when a magnetic domain is transferred out of its virtual position in a memory storage loop at the second point on that loop by the second one-way gate, a new magnetic domain may be transferred into the same virtual position by the first one-way gate when the virtual position next reaches the first point on that loop.

6 Claims, 5 Drawing Figures

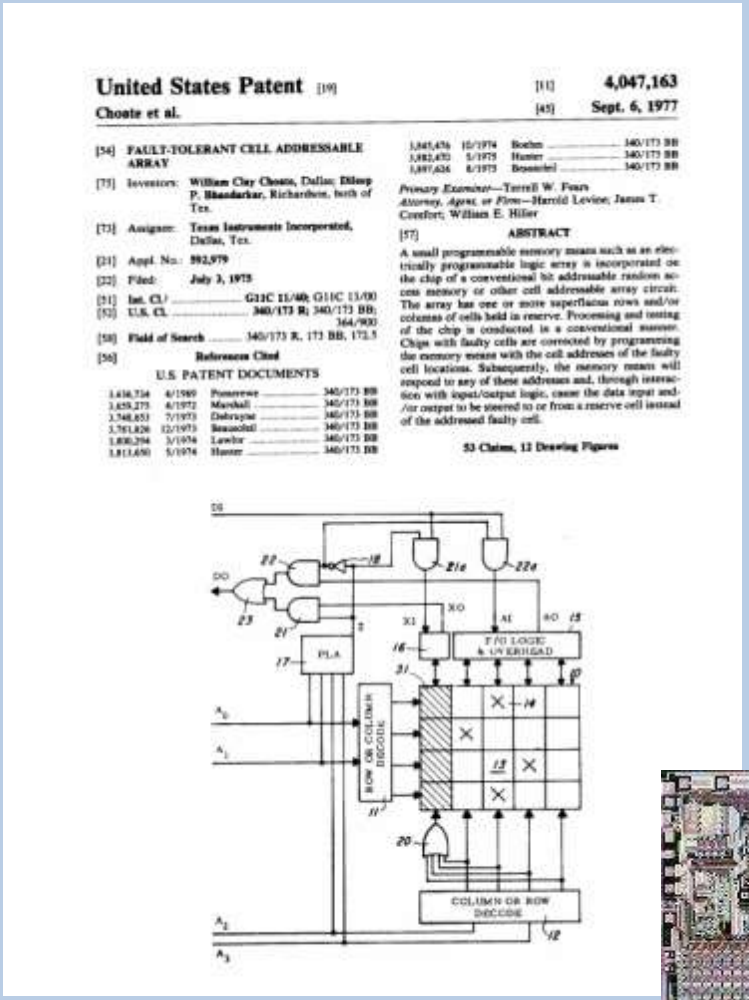

United States Patent [19]
Choate et al.
[11] 4,047,163
[45] Sept. 6, 1977

[54] FAULT-TOLERANT CELL ADDRESSABLE ARRAY

[75] Inventors: William Clay Choate, Dallas; Dileep P. Bhandarkar, Richardson, both of Tex.

[73] Assignee: Texas Instruments Incorporated, Dallas, Tex.

[21] Appl. No.: 592,979

[22] Filed: July 3, 1975

[51] Int. Cl.² G11C 11/40; G11C 13/00
[52] U.S. Cl. ..... 340/173 R; 340/173 BB; 364/900
[58] Field of Search ..... 340/173 R, 173 BB, 172.5

[56] References Cited
U.S. PATENT DOCUMENTS

Primary Examiner—Terrell W. Fears
Attorney, Agent, or Firm—Harold Levine; James T. Comfort; William E. Hiller

[57] ABSTRACT

A small programmable memory means such as an electrically programmable logic array is incorporated on the chip of a conventional bit addressable random access memory or other cell addressable array circuit. The array has one or more superfluous rows and/or columns of cells held in reserve. Processing and testing of the chip is conducted in a conventional manner. Chips with faulty cells are corrected by programming the memory means with the cell addresses of the faulty cell locations. Subsequently, the memory means will respond to any of these addresses and, through interaction with input/output logic, cause the data input and/or output to be steered to or from a reserve cell instead of the addressed faulty cell.

53 Claims, 12 Drawing Figures

TMS 9900

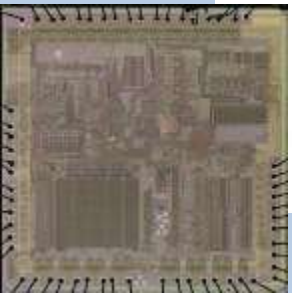

TMS 1000

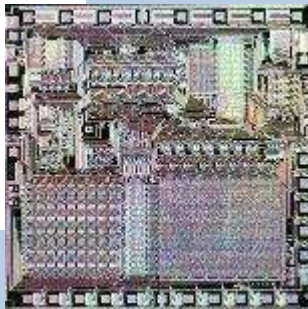

IBM 370/168 – circa 1974

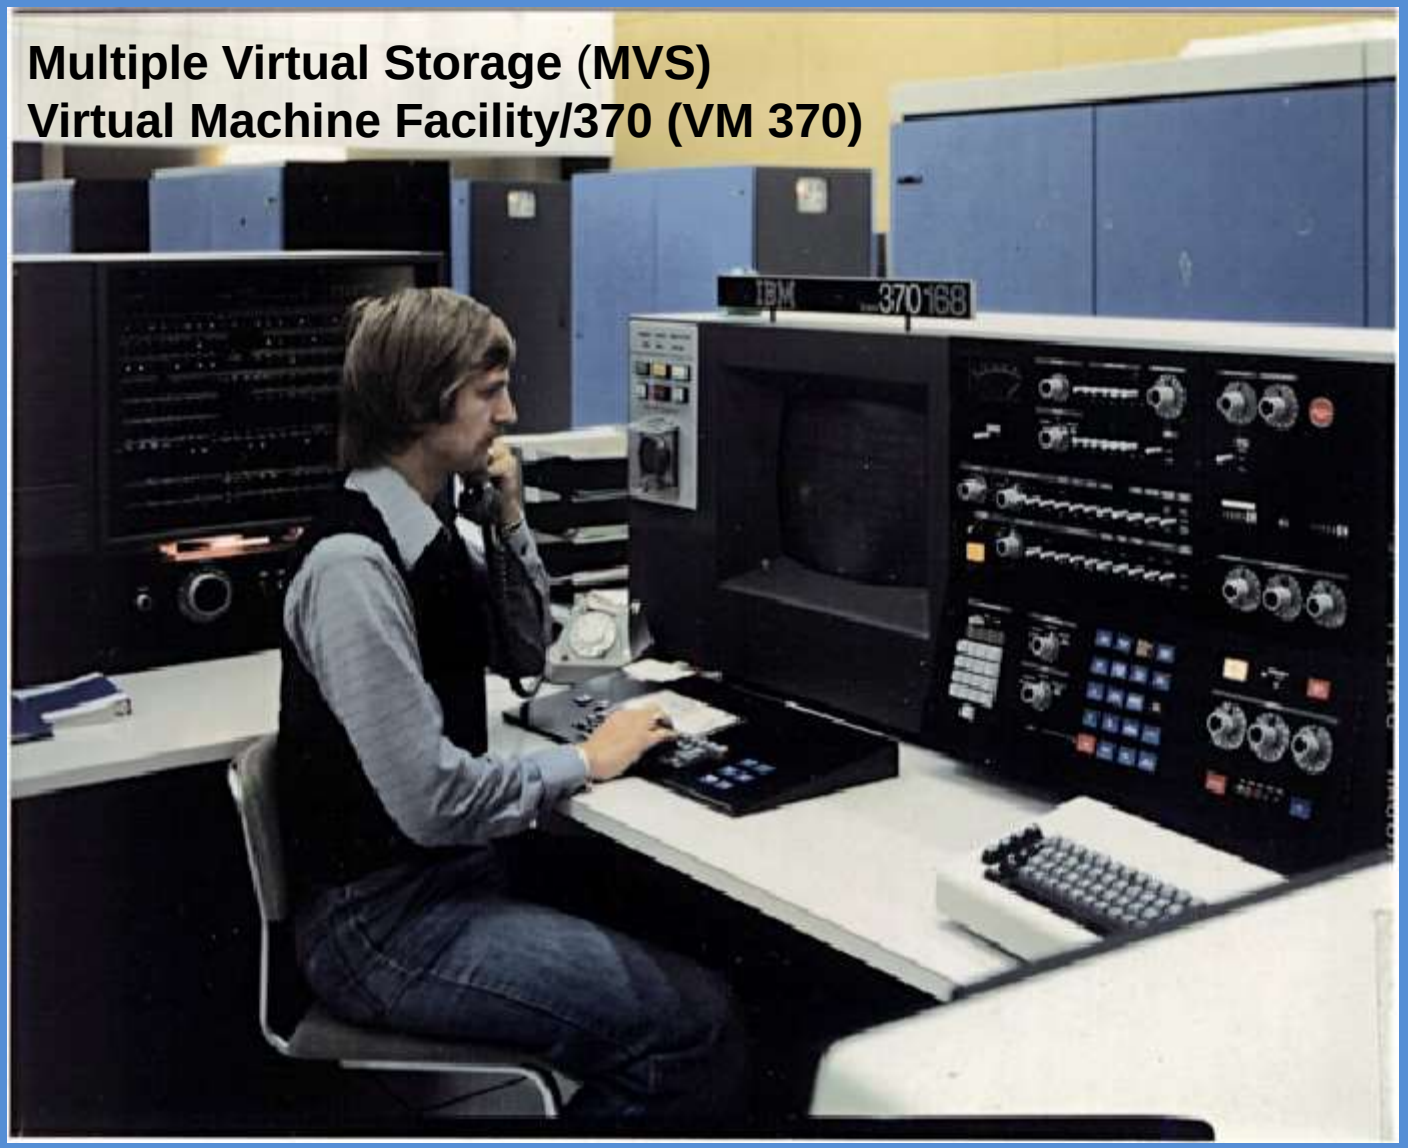

"I think there is a world market for maybe five computers", Thomas J. Watson (1943)


1978 – 1995
17.5 Year Odyssey
at
Digital Equipment Corporation
http://research.microsoft.com/en-us/um/people/gbell/Digital/timeline/1978.htm

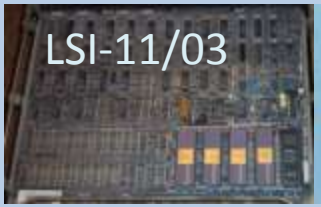

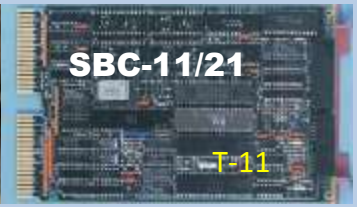

DEC PDP-11

PDP-11/20: original, non-microprogrammed processor

microprogrammed successor to the PDP-11/20

http://www.psych.usyd.edu.au/pdp-11/models.html

1977: VAX-11/780 – STAR is Born!

VAX Family: 1977 - 1992

Evolution of VAX Architecture:

- **Floating Point Extensions**
- **MicroVAX Subset**
- **VAXvector Extensions**

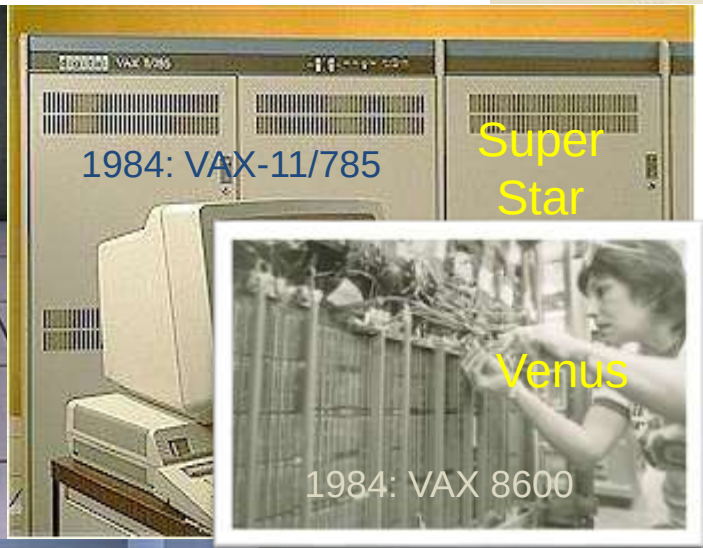

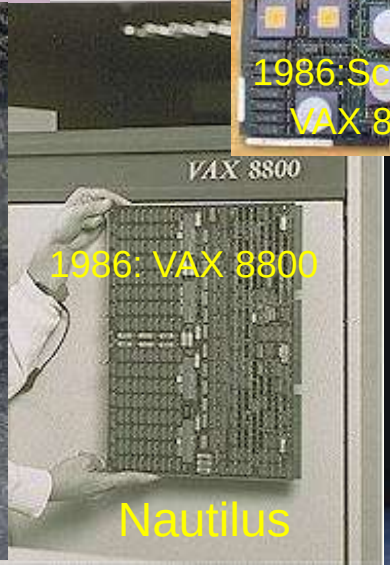

1985: MicroVAX-II (Subset ISA)

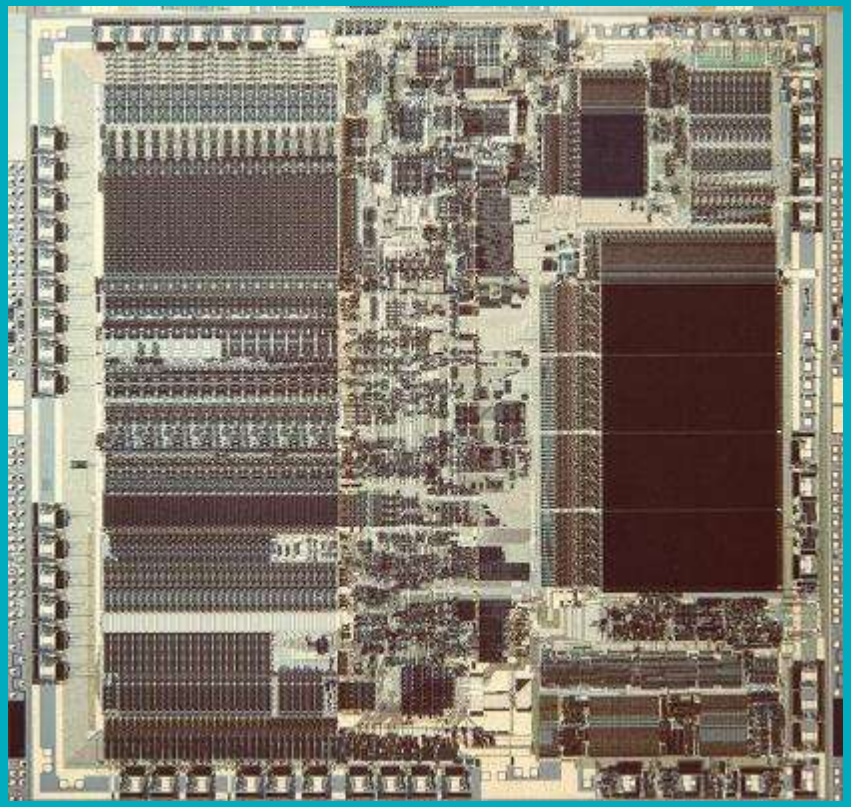

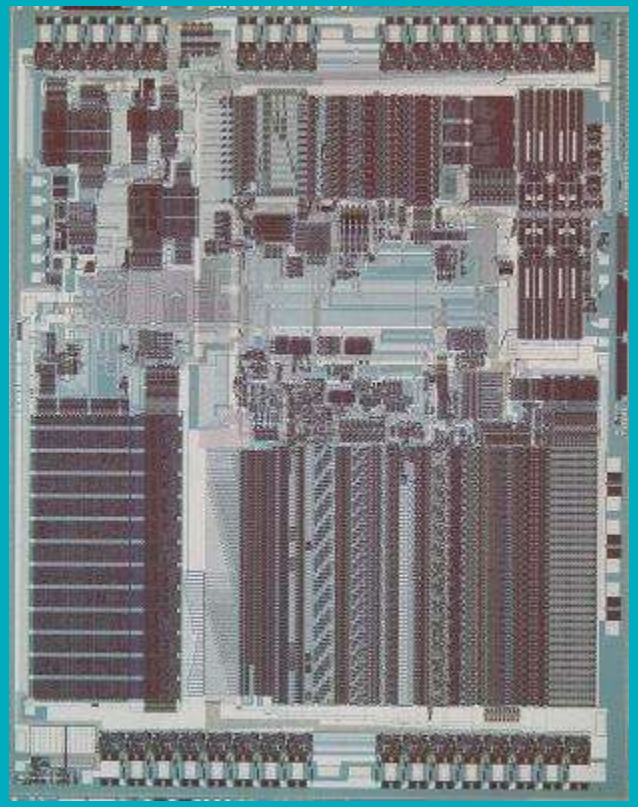

The **MicroVAX- I** (*Seahorse*), introduced in October 1984, was the first VAX computer to use VLSI Technology.

1988-1992: VAX 6000 Series

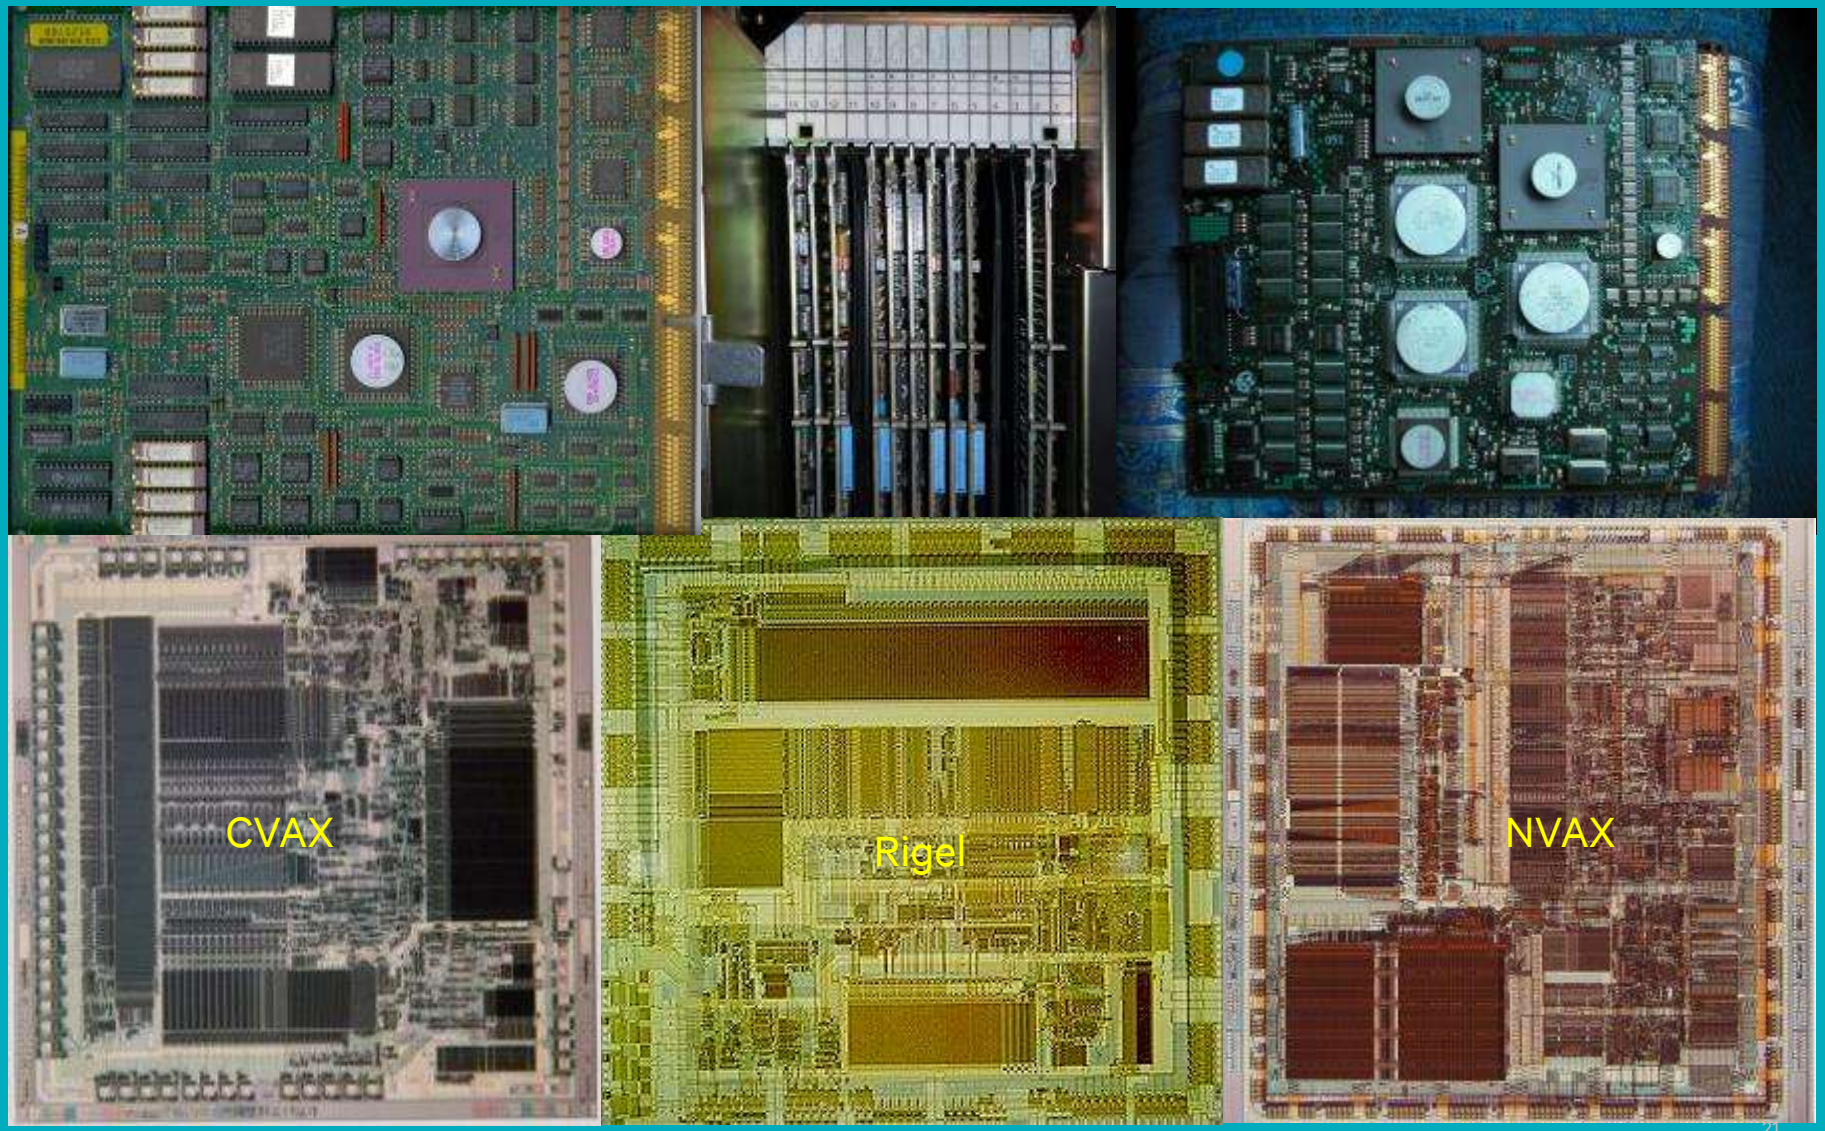

when you care enough to steal the very best

1989: VAX 9000 - The Age of Aquarius

The Beginning of the End of VAX

RISC vs CISC WARS

- Sun SPARC
- MIPS R2000, R3000, R4000, R6000, R10000
- PA-RISC
- IBM Power and Power PC
- DEC Alpha 21064, 21164, 21264

microPRISM

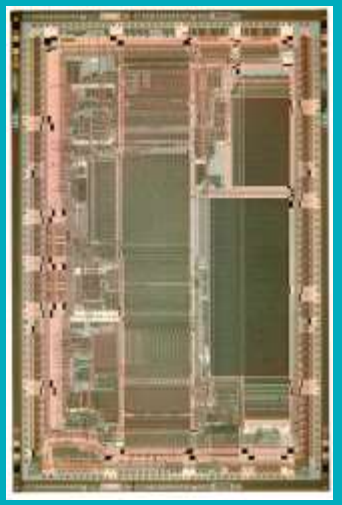

In 1987, the introduction of RISC processors based on Sun's SPARC architecture spawned the now famous RISC vs CISC debates. DEC cancelled PRISM in 1988. RISC processors from MIPS, IBM (Power, Power PC), and HP (PA-RISC) started to gain market share. This forced Digital to adopt MIPS processors in 1989, and later introduce Alpha AXP (Almost eXactly PRISM!) in 1992.

Alpha 21064

High Performance Issue Oriented Architecture

D. Bhandarkar, D. Orbits, R. Witek, W. Cardoza, D. Cutler†

Digital Equipment Corporation

RISC vs CISC Debate

- VAX was king of CISC
  - More than 300 instructions of variable length
  - Compact code size
  - Hard to decode quickly
  - Low Freq, Short Path Length, Complex Design
- Iron Law of Performance:
  - Speed = IPC * freq /Path Length
- RISC championed by SPARC and MIPS
  - Simpler instruction format but longer path length
  - Higher frequency (Brainiacs vs Speed Demons)
- RISC was “better” for in order designs
- Out of order microarchitectures leveled the playing field
- Semiconductor Technology and Volume Economics matter!
- PC Volumes and Pentium Pro design changed the industry

The difference between theory and practice is always greater in practice than it is in theory!

Performance from Architecture: Comparing a RISC and a CISC with Similar Hardware Organization

Dileep Bhandarkar
Digital Equipment Corp.
146 Main Street (ML05-2/G1)
Maynard, MA 01754

Douglas W. Clark*
Allen Computation Lab
Harvard University
Cambridge, MA 02138

Abstract

1 Introduction

310

RISC versus CISC: A Tale of Two Chips

Dileep Bhandarkar
Intel Corporation
Santa Clara, California, USA

Abstract

Introduction

Chip Overview

Table 1 Chip Comparison

	Alpha 21164	Pentium® Pro Processor
Architecture	Alpha	IA-32
Clock Speed	300 MHz	150 MHz
Issue Rate	Four	Three
Function Units	four	five
Out of order issue	no	yes
Rename Registers	none	40
On-chip Cache	8 KB data 8KB inst 96 KB Level 2	8 KB data 8KB inst
Off chip cache	4 MB	256 KB
Branch History	2048 entries,	512 entries,
Table	2-bit history	4-bit history
Transistors		
Logic	1.8 million	4.5 million
Total	9.3 million	5.5 million
VLSI Process	CMOS	BiCMOS
Min. Geometry	0.5 μ	0.6 μ
Metal Layers	4	4
Die Size	298 mm²	306 mm²
Package	499 pin PGA	387 pin PGA
Power	50 W	28 W incl cache
First Silicon	Feb 94	4Q 94
Volume Parts	3Q 95	4Q 95
SPECint92/95	341/7.83	365/6.88
SPECfp92/95	513/12.4	256/5.42
SYSmark/NT	523	447

1991: ACE Initiative

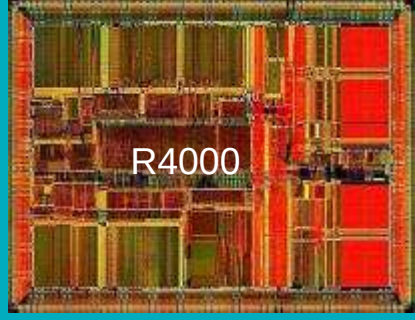

- The Advanced Computing Environment (ACE) was defined by an industry consortium in the early 1990s to be the next generation commodity computing platform, the successor to personal computers based on Intel's 32-bit x86 instruction set architecture.
- The consortium was announced on the 9th of April 1991 by MIPS Computer Systems, Digital Equipment Corporation, Compaq, Microsoft, and the Santa Cruz Operation.
- **At the time it was widely believed that RISC-based systems would maintain a price/performance advantage over the x86 systems.**
- The environment standardized on the MIPS architecture and two operating systems: SCO UNIX with Open Desktop and what would become Windows NT (originally named OS/2 3.0).
- The Advanced RISC Computing (ARC) document was produced to define hardware and firmware specifications for the platform.
- When the initiative started, MIPS R3000 RISC based systems had substantial performance advantage over Intel 80486 and original 60 MHz Pentium chips .
- **MIPS R4000 schedule and performance slipped and Intel updated the Pentium design to 90 MHz in the next semiconductor process generation and the MIPS performance advantage slipped away.**

Strategy without Execution is Doomed!

Alpha was Too Little Too Late!

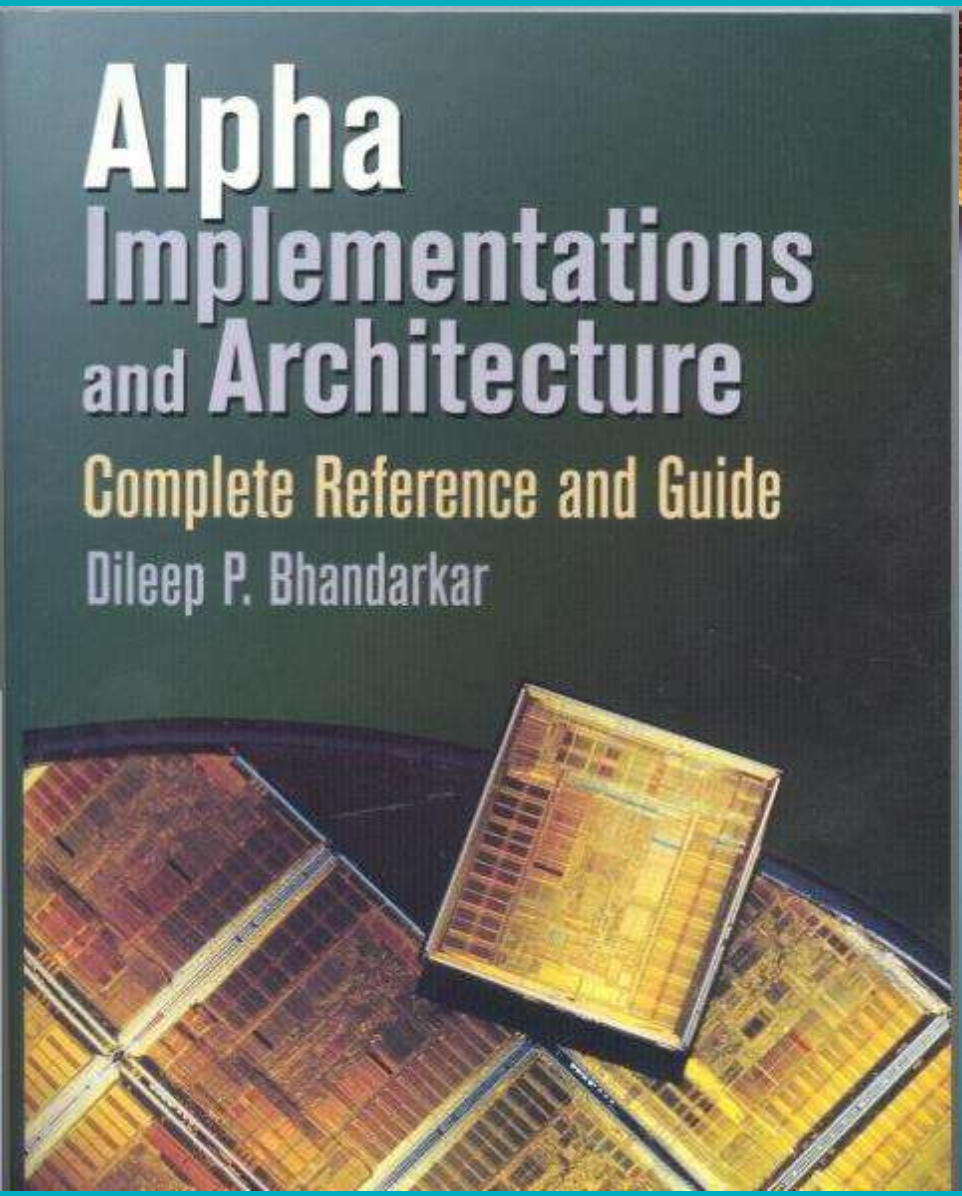

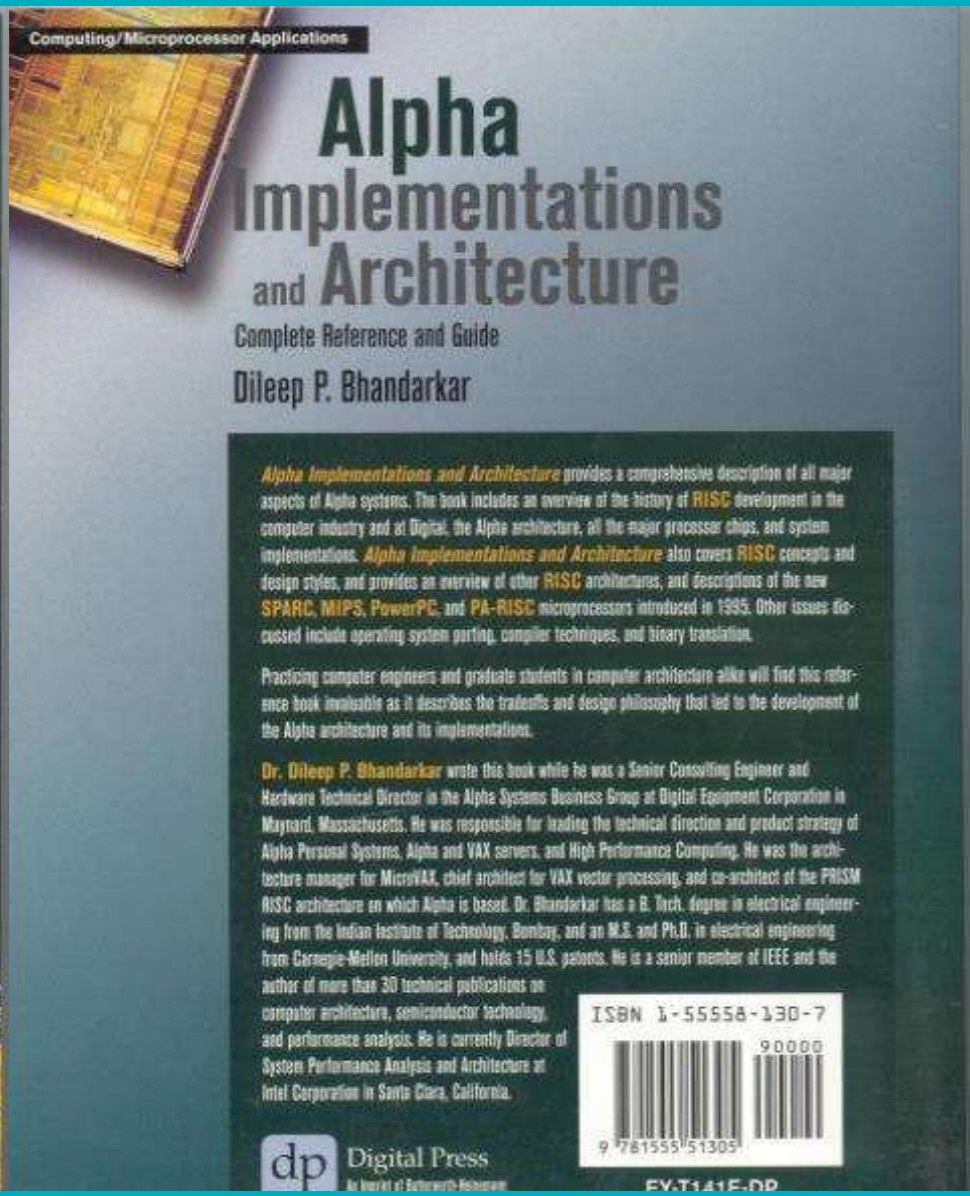

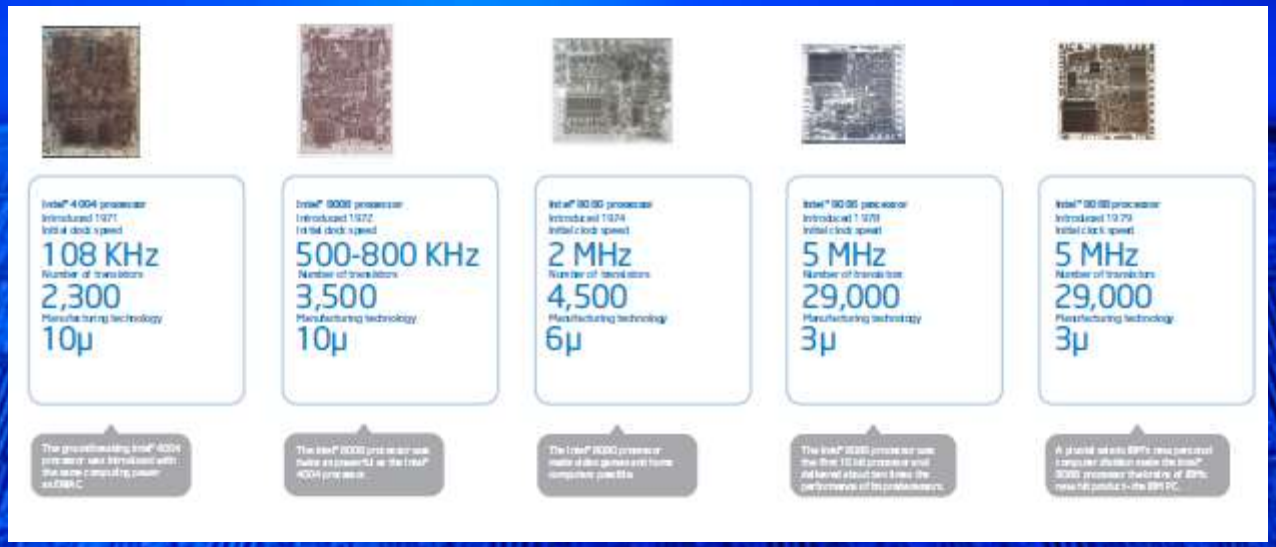

Looking at Intel from the Outside

1974: 8080 Microprocessor

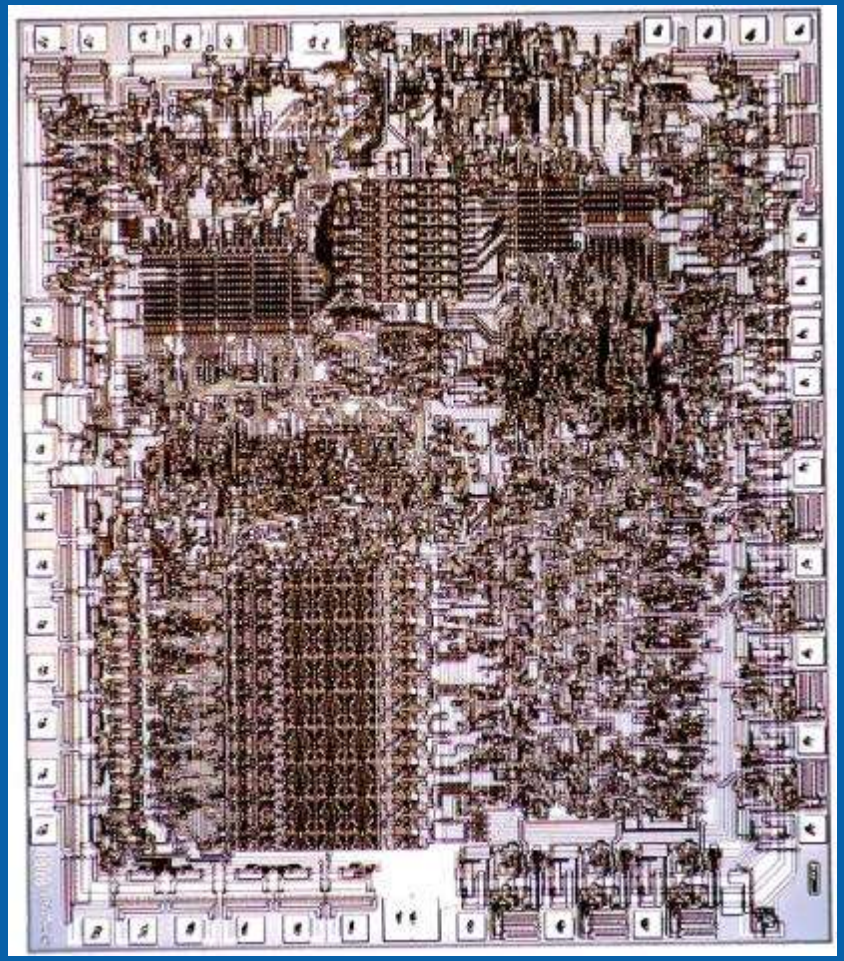

- The 8080 became the brain of the first personal computer--the Altair, allegedly named for a destination of the Starship *Enterprise* from the *Star Trek* television show. Computer hobbyists could purchase a kit for the Altair for $395.
- Within months, it sold tens of thousands, creating the first PC back orders in history
- 2 MHz
- 4500 transistors
- 6 µm

1978-79: 8086-8088 Microprocessor

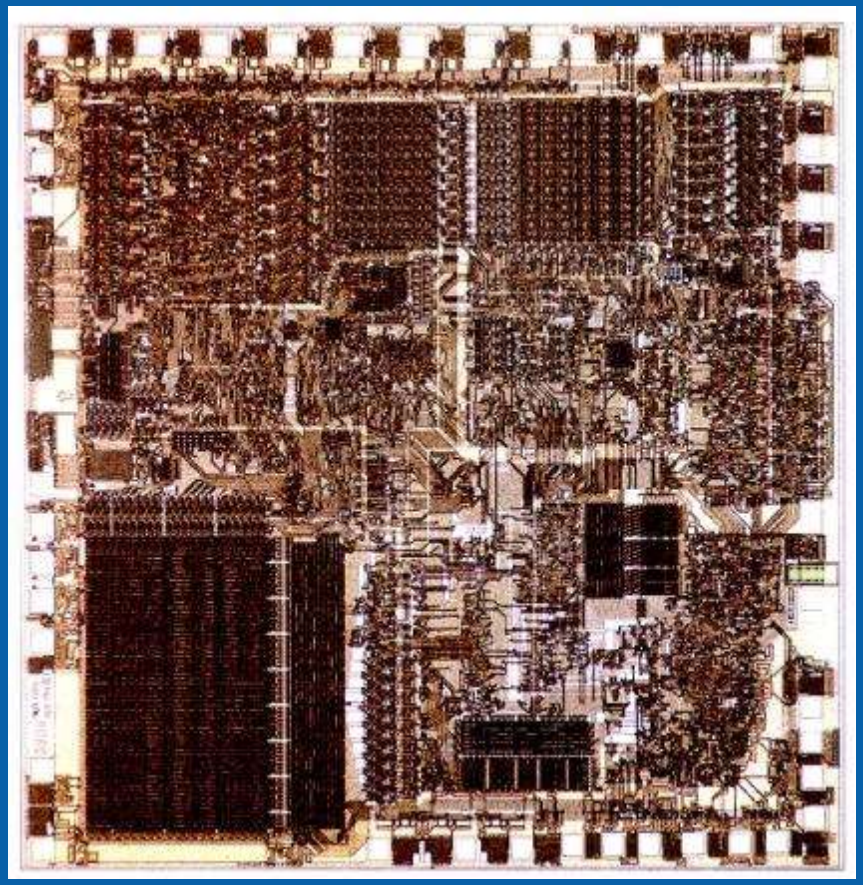

- A pivotal sale to IBM's new personal computer division made the 8088 the brain of IBM's new hit product--the IBM PC.
- The 8088's success propelled Intel into the ranks of the *Fortune 500*, and *Fortune* magazine named the company one of the "Business Triumphs of the Seventies."
- 5 MHz
- 29,000 transistors
- 3 µm

1981: First IBM PC

The IBM Personal Computer ("PC")

- PC-DOS Operating System
- Microsoft BASIC programming language, which was built-in and included with every PC.
- Typical system for home use with a memory of 64K bytes, a single diskette drive and its own display, was priced around $3,000.
- An expanded system for business with color graphics, two diskette drives, and a printer cost about $4,500.

"There is no reason anyone would want a computer in their home." Ken Olsen, president Digital Equipment Corp (1977)

1979: Motorola 68000

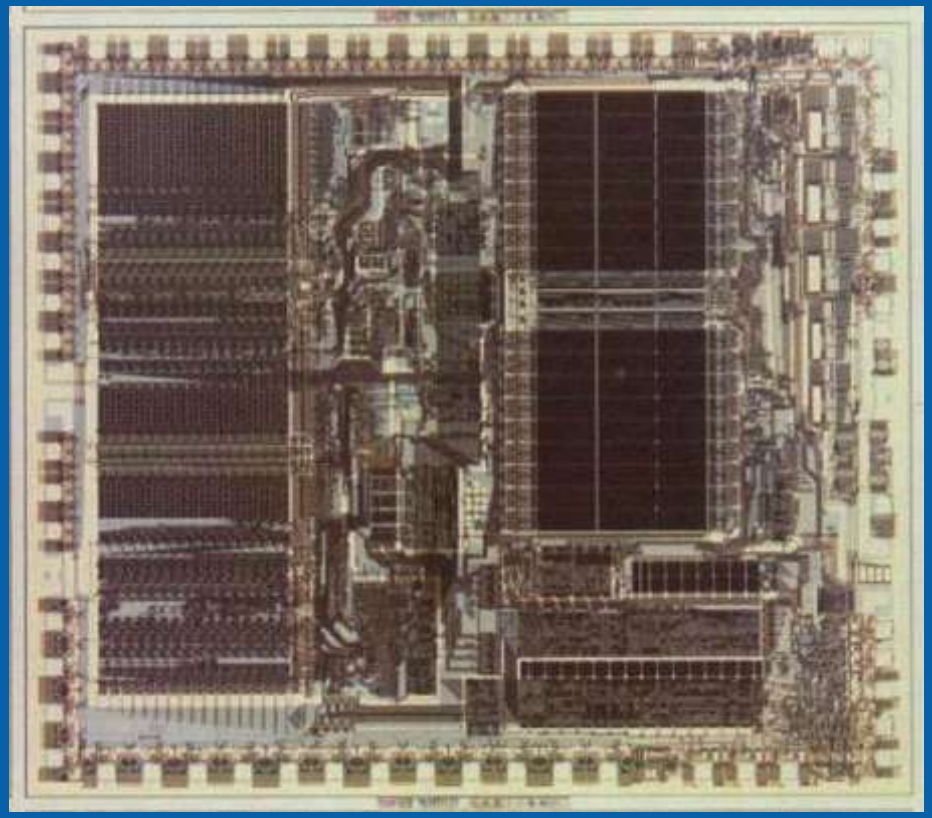

1984: Apple Macintosh

The 68000 became the dominant CPU for Unix-based workstations from Sun and Apollo

It was also used for personal computers such as the Apple Lisa, Macintosh, Amiga, and Atari ST

1985: Intel386™ Microprocessor

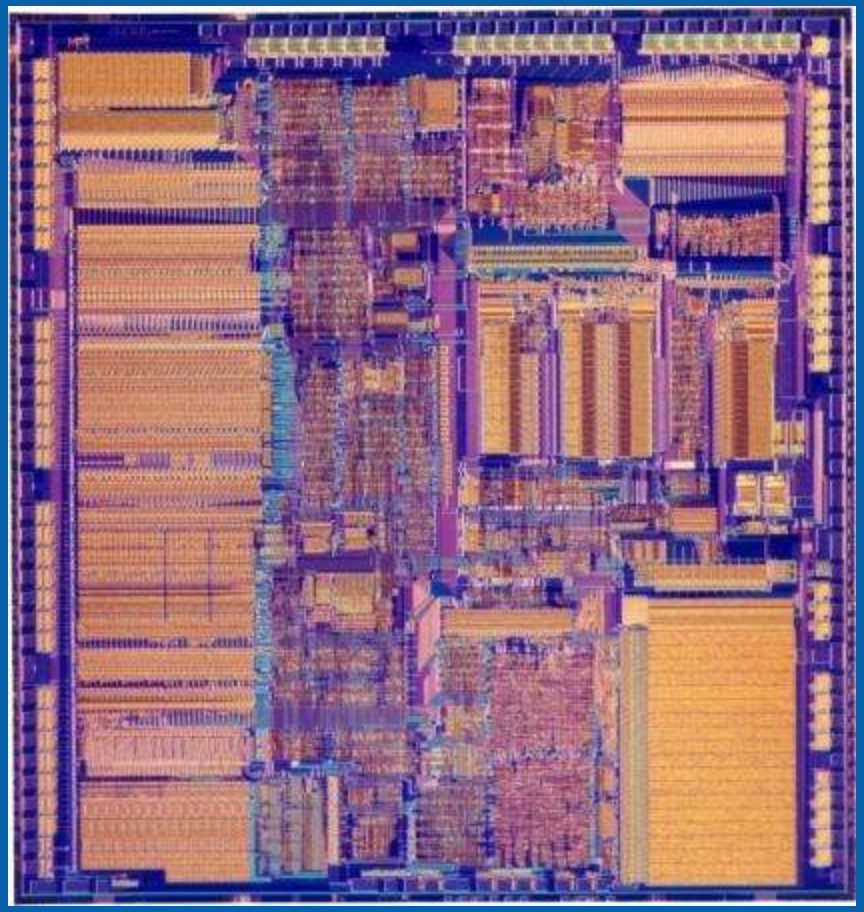

- The Intel386™ microprocessor featured 275,000 transistors--more than 100 times as many as the original 4004. It was a Intel's first 32-bit chip.
- The 80386 included a paging translation unit, which made it much easier to implement operating systems that used virtual memory.
- 16 MHz
- 1.5µm

1989: Intel486™ DX CPU Microprocessor

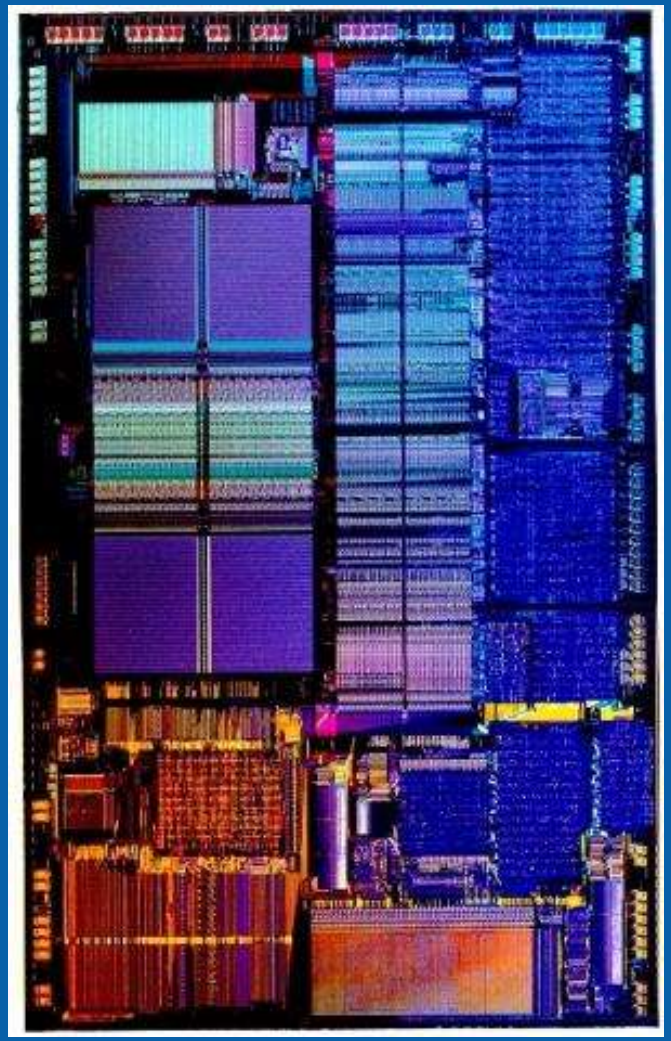

- The Intel486™ processor was the first to offer a "large" 8KB unified instruction and data on-chip cache and an integrated floating-point unit.
- Due to the tight pipelining, sequences of simple instructions (such as ALU reg, reg and ALU reg, im) could sustain a single clock cycle throughput (one instruction completed every clock).
- 25 MHz
- 1.2 M transistors
- 1 µm

1993: Intel® Pentium® Processor

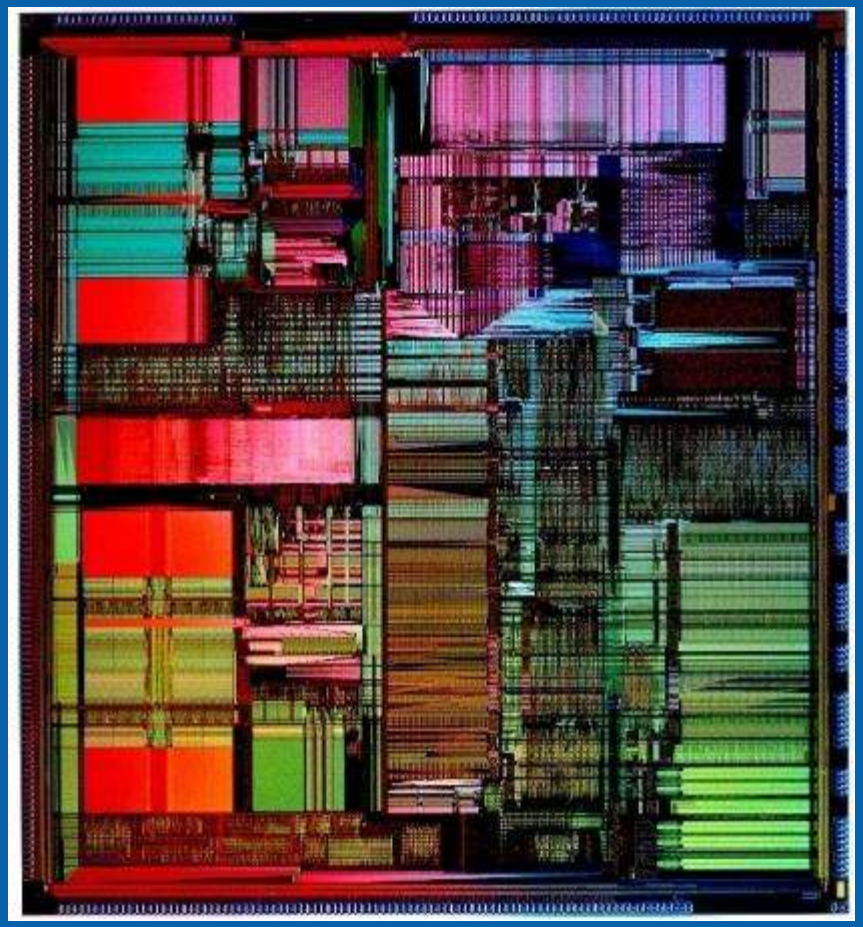

- The Intel Pentium® processor was the first superscalar x86 microarchitecture. It included dual integer pipelines, a faster floating-point unit, wider data bus, separate instruction and data caches
- Famous for the FDIV bug!
- 22 March 1993
- 66 MHz
- 3.1 M transistors
- 0.8 µm

Start of the sub-micron era!

1995: Intel® Pentium® Pro Processor

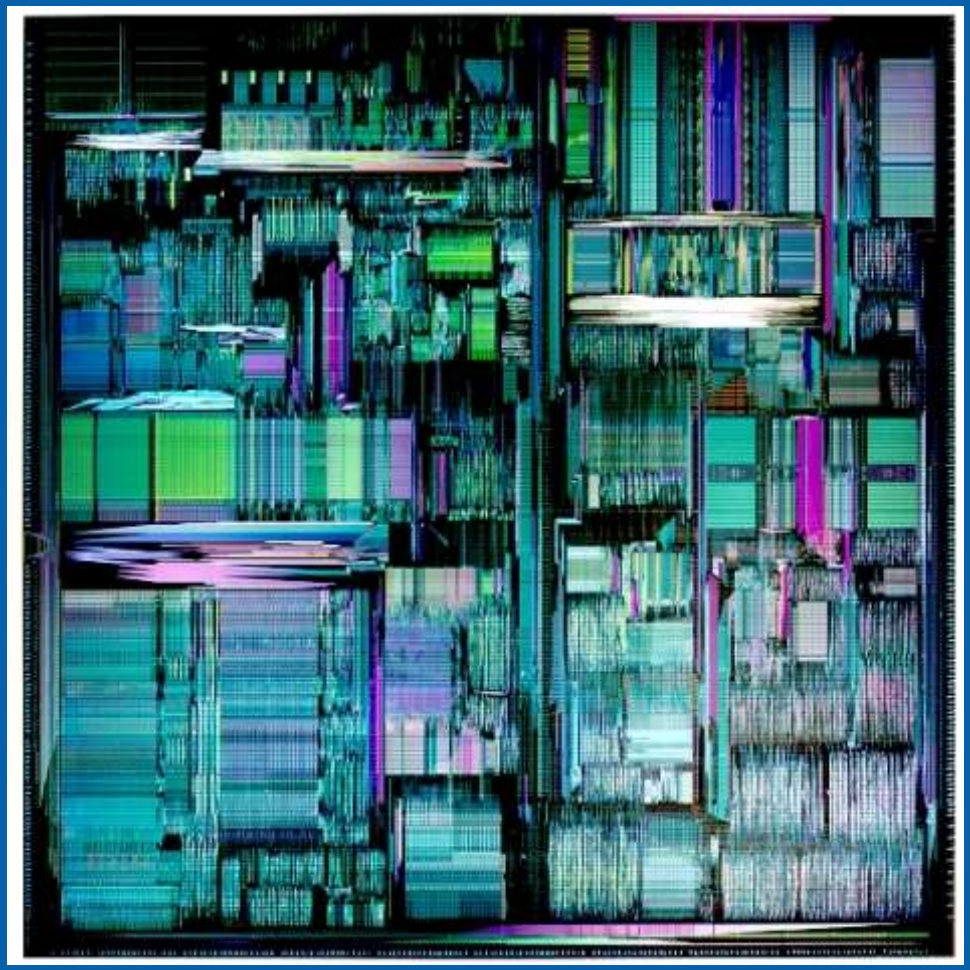

- Intel® Pentium® Pro processor was designed to fuel 32-bit server and workstation applications. Each processor was packaged together with a second L2 cache memory chip on the back-side bus.
- 5.5 million transistors.
- 1 November 1995
- 200 MHz
- 0.35µm
- 1st x86 to implement out of order execution
- Front side bus with split transactions
- The P6 micro-architecture lasted 3 generations from the Pentium Pro to Pentium III
- The Pentium Pro processor slightly outperformed the fastest RISC microprocessors on integer benchmarks, but floating-point performance was significantly lower

The RISC Killer!

June 1995: Inside Intel

If You Can't Beat Them Join Them!

1997-98: Intel® Pentium® II Processor

Klamath
Deschutes

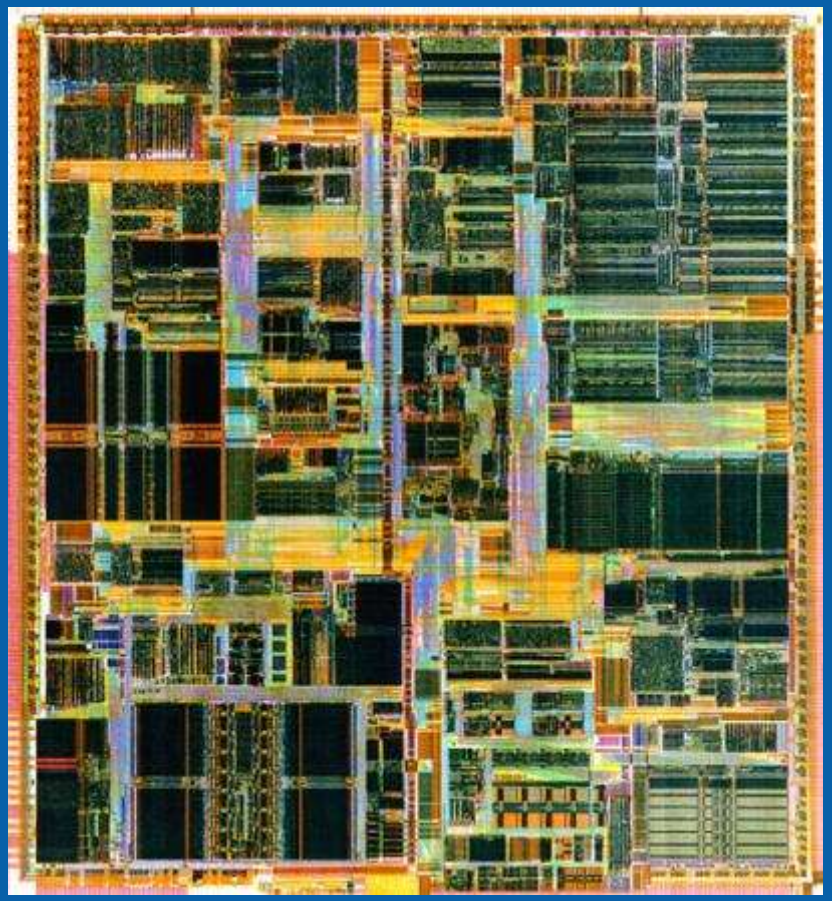

- The 7.5 million-transistor 0.35 µm Pentium II processor was introduced with 512 KB L2 cache in external chips on the CPU module clocked at half the CPU's 300 MHz frequency in a "Slot 1" SECC module.
- **1998: Intel Pentium II Xeon** processors (0.25 µm Deschutes) were launched with a full-speed custom 512 KB, 1 MB, or 2 MB L2 cache using a larger Slot 2 to meet the performance requirements of mid-range and higher servers and workstations

Mendocino

1998: Intel® Celeron® Processor

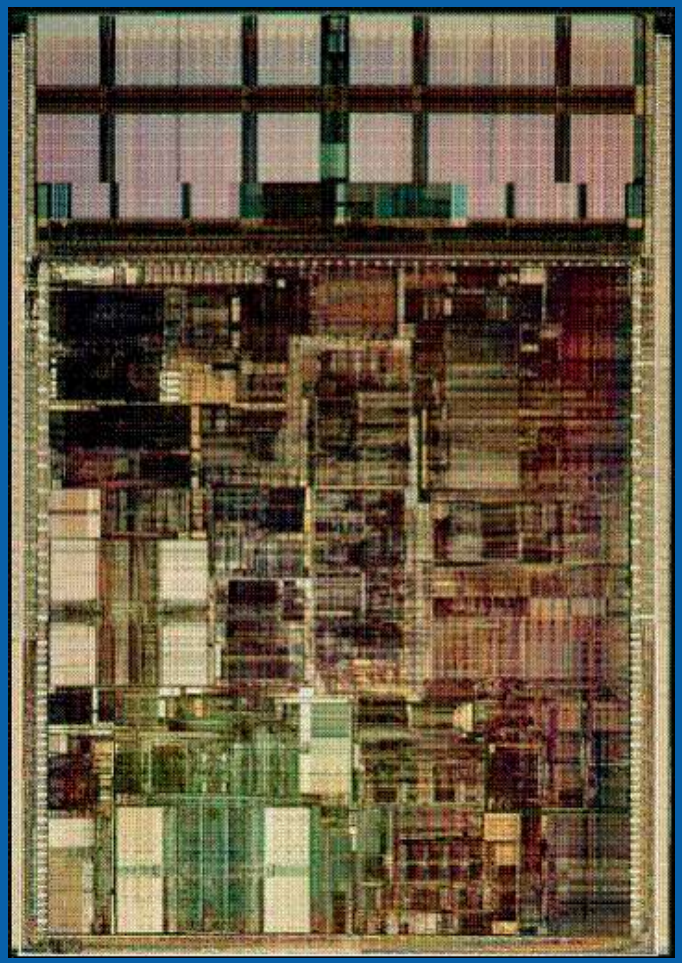

- The Intel® Celeron® processors were designed for the sub $1000 Value PC market segment.
- The first Celeron processor (Covington) in April 1998 was just a 266 MHz Pentium II without a L2 cache
- Mendocino: First x86 with integrated L2 cache -128 KB
- 19M transistors
- 300 MHz
- 0.25µm
- 24 August 1998

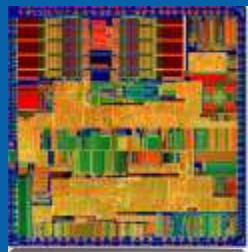

Intel's Response to Cyrix 6x86 (M1)

1999: AMD Athlon

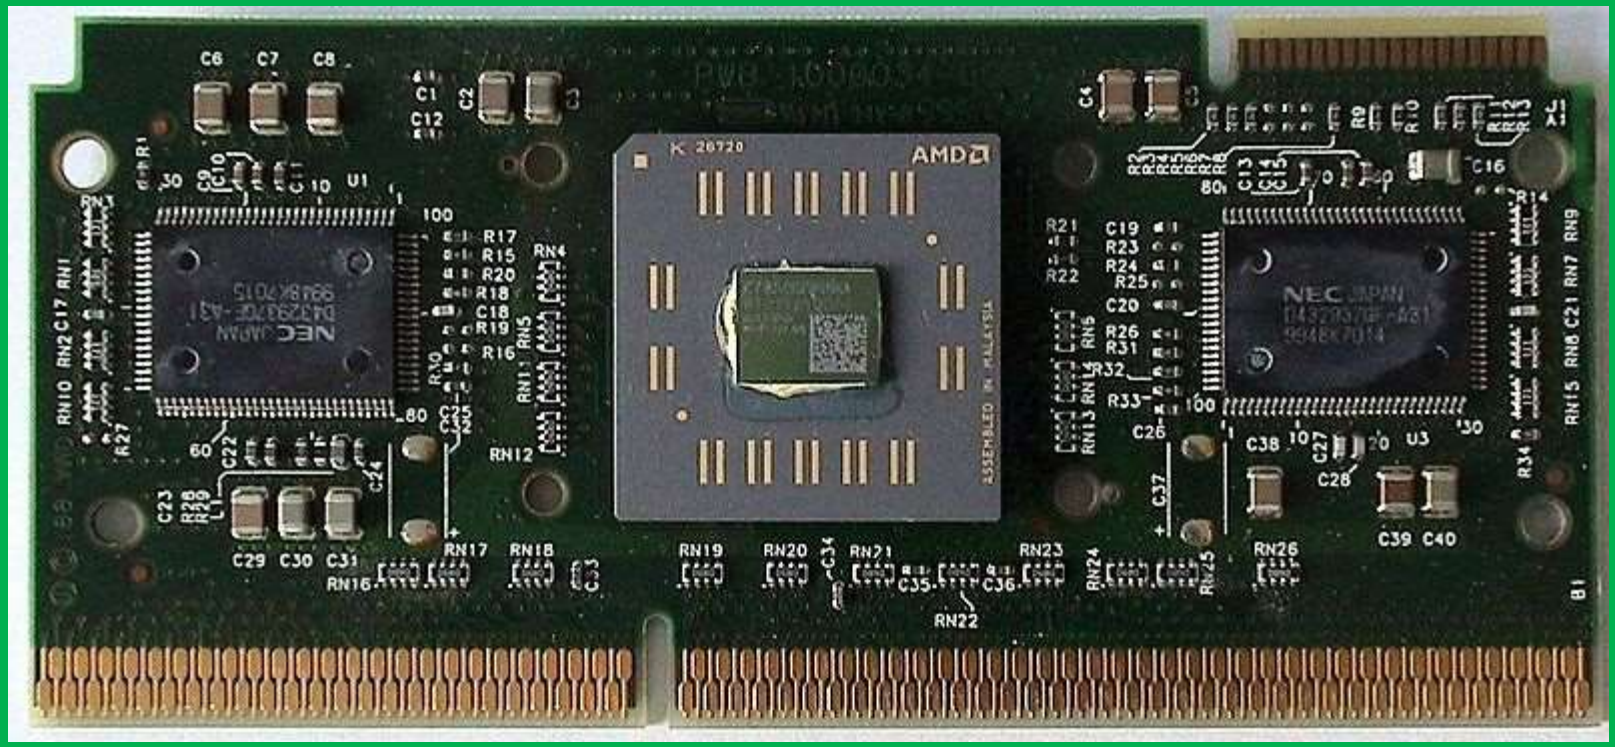

Won the Race to 1 GHz

Oct 2009: AMD Hammers Intel with AMD64

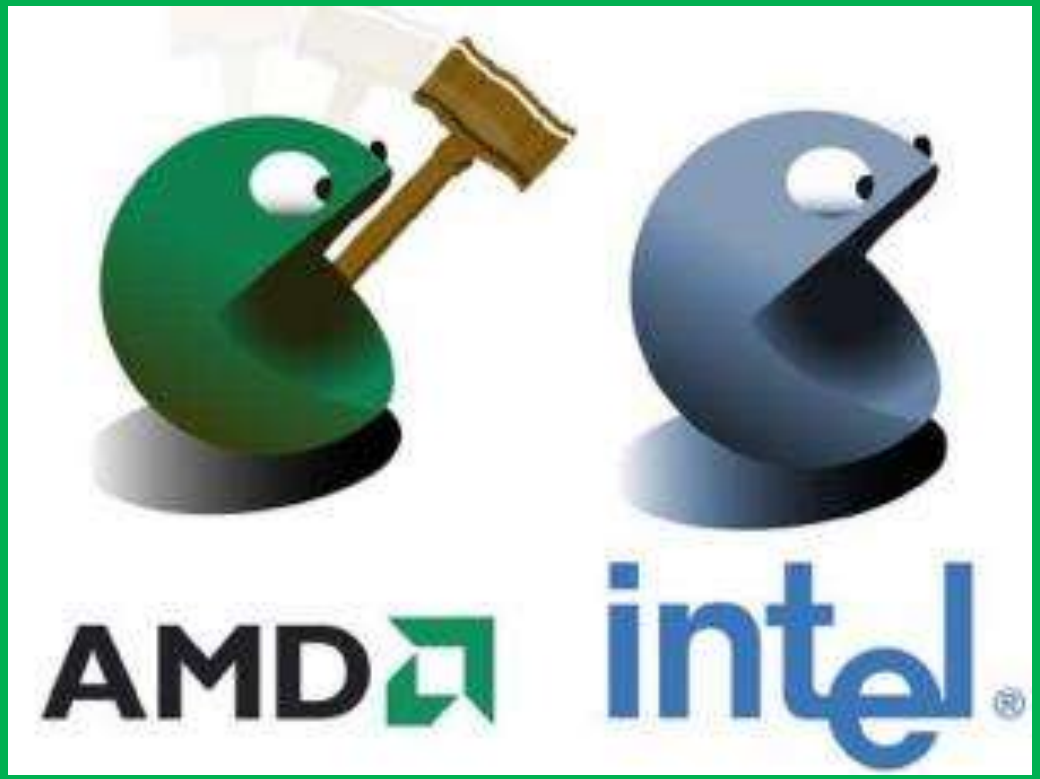

5 Oct 2009, SAN JOSE, California--Advanced Micro Devices today is detailing a new 64-bit chip that will compete against Intel's Itanium processor. The chip will be an extension of the current Intel-compatible chip design, or so-called *x*86 architecture, said Fred Weber, vice president of engineering at AMD's computation products group, at a processor industry conference here today. Intel's next-generation design, Itanium, will be a wholly new architecture.

Coppermine

1999: Intel® Pentium® III Processor – 0.18µm

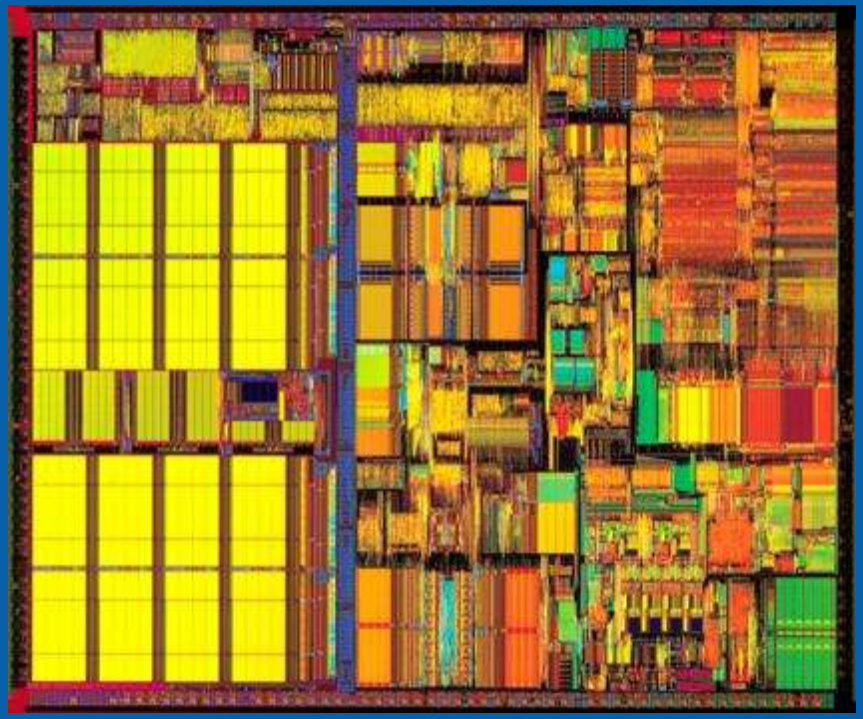

- 25 Oct 1999
- Integrated 256KB L2 cache
- 733 MHz
- 28 M transistors
- 1st Intel microprocessor to hit 1 GHz on 8-Mar-2000, a few days after AMD Athlon!

Willamette

2000: Intel® Pentium® 4 Processor – 0.18µm

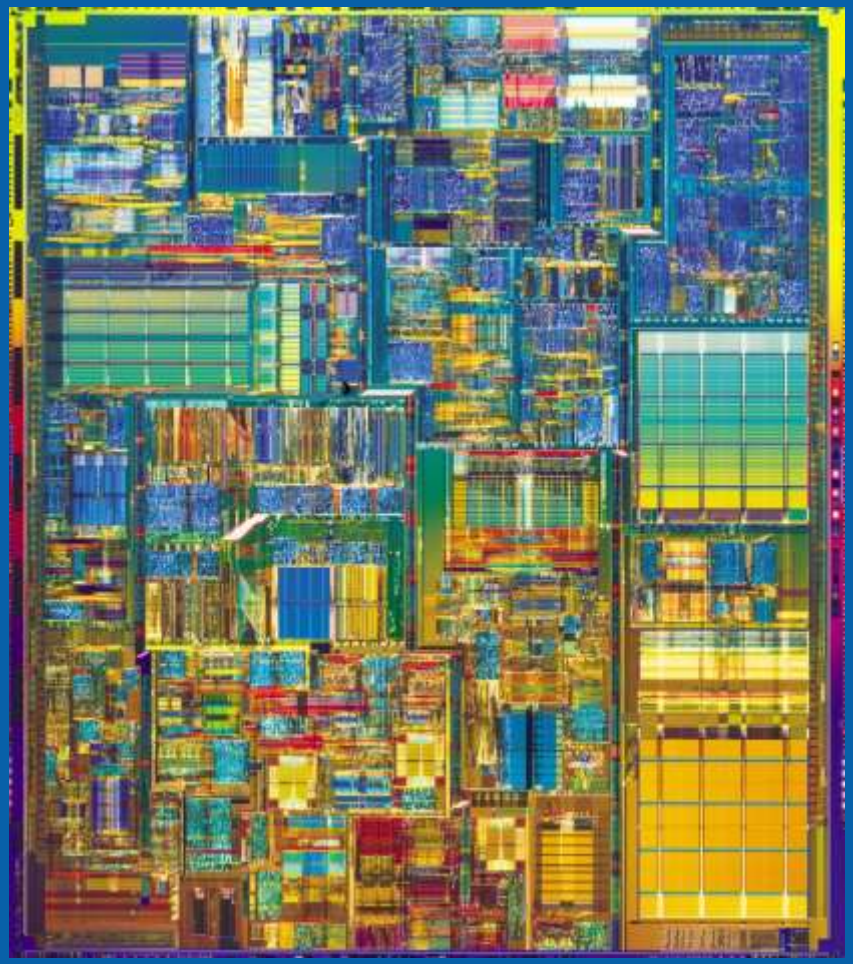

- The Intel® Pentium® 4 processor's initial speed was 1.5 GigaHertz.
- 20 Nov 2000
- 256K integrated L2 cache
- Double clocked "Fireball" inner core
- Deep 20 stage pipeline
- 100 MHz quad pumped bus
- 42 M transistors
- Hit 2 GHz on 27 Aug 2001
- ~55 Watts
- No Mobile Pentium 4!

High Frequency, but Power was High too!

Merced

2001: Intel® Itanium™ Processor

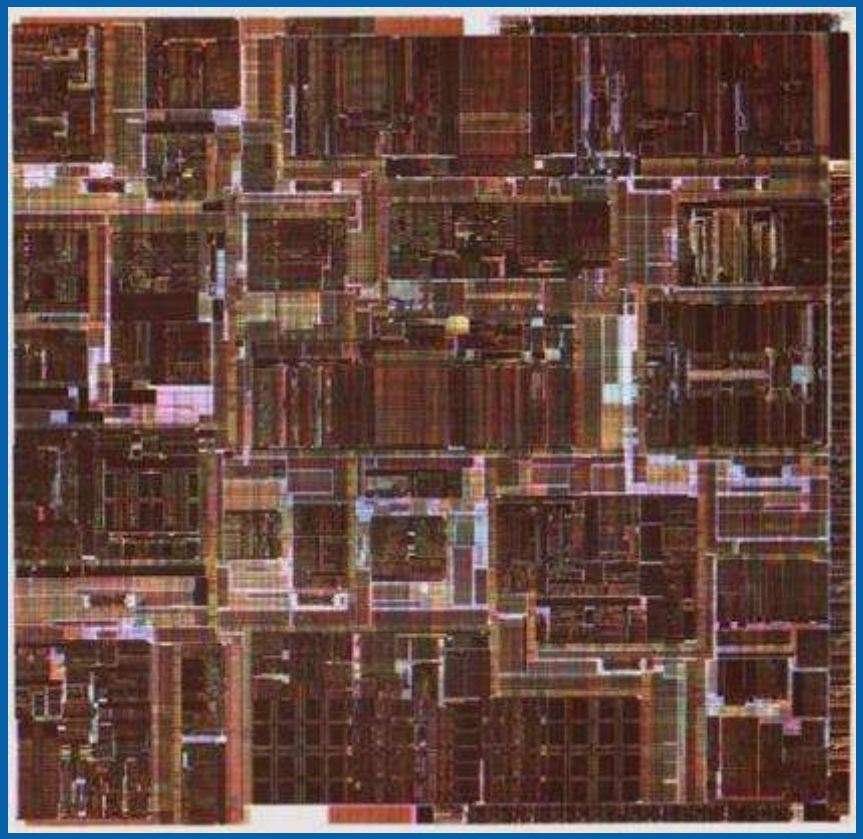

- The Itanium™ processor is the first in a family of 64-bit products from Intel. Designed for high-end, enterprise-class servers and workstations, the processor was built from the ground up with an entirely new architecture based on Intel's Explicitly Parallel Instruction Computing (EPIC) design technology.
- Based on HP's VLIW project
- May 2001
- 800 MHz
- 25M transistors
- 0.18µm
- 4 MB External Level 3 Cache
- Intel's EPIC Blunder!

IT AIN'T PENTIUM !!!

Northwood

2001: Intel® Pentium® 4 Processor – 0.13µm

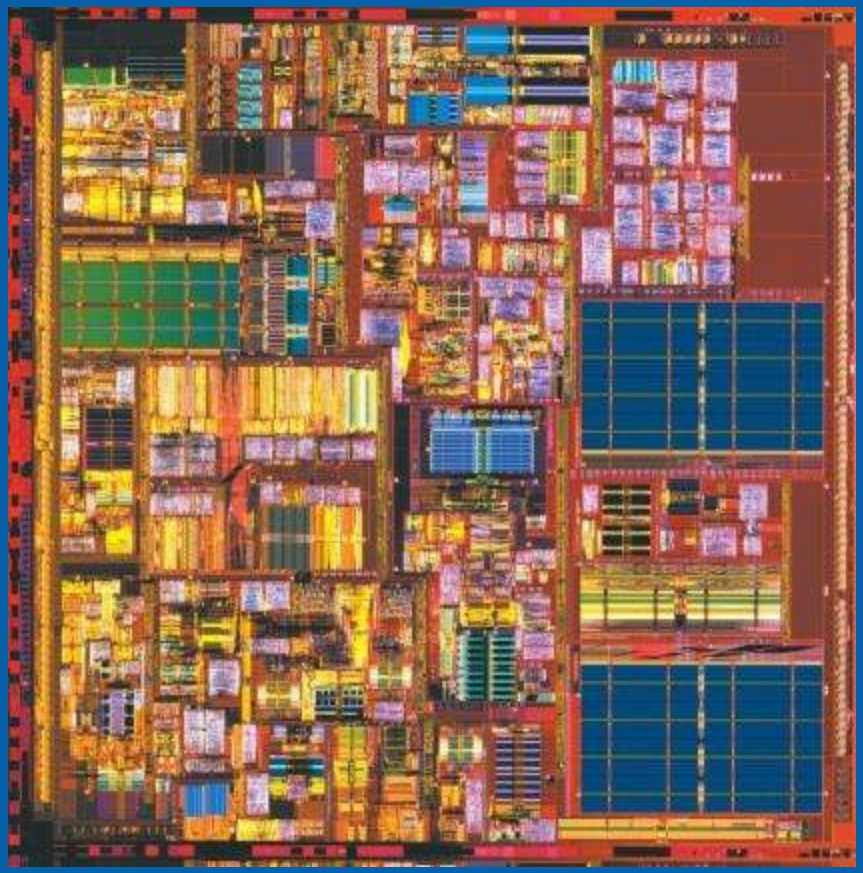

- 27 August 2001
- 55 million transistors
- 2 GHz
- 512KB L2 cache
- In 2002 Intel released a Xeon branded CPU, codenamed "Prestonia" with Intel's Hyper-Threading Technology
- 14 Nov 2002: 3.06 GHz
- 23 June 2003: 3.2 GHz

Simultaneous Multi Threading Introduced to x86 Processors

2003: AMD Opteron – First 64 bit x86

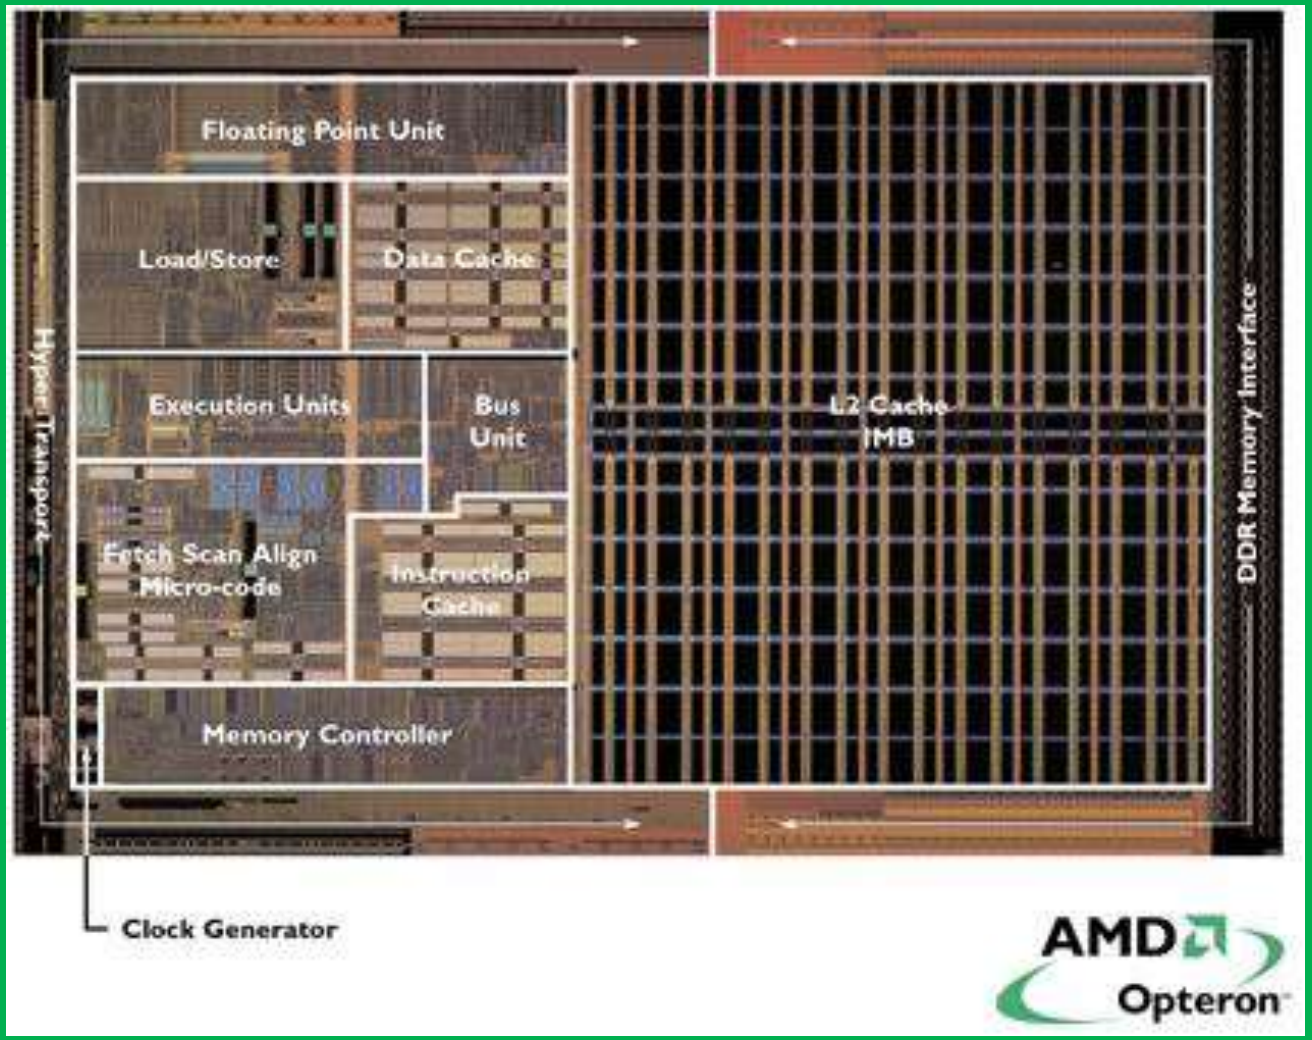

First x86 processor with 64 bits and Integrated Memory Controller

2003: Intel® Pentium® M Processor

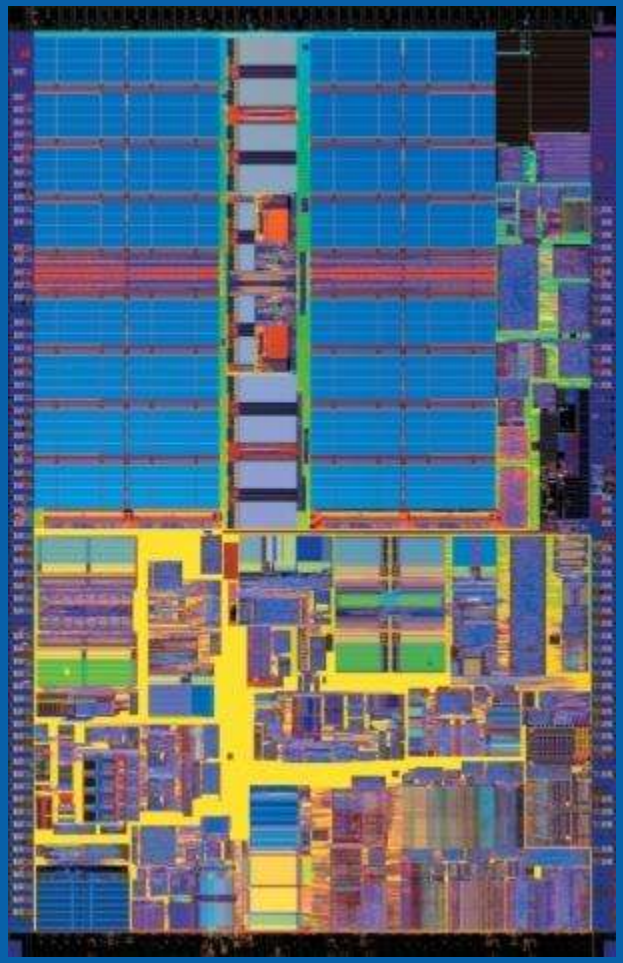

- The first Intel® Pentium® M processor, the Intel® 855 chipset family, and the Intel® PRO/Wireless 2100 network connection were the three components of Intel® Centrino™ mobile technology, with built-in wireless LAN capability and breakthrough mobile performance. It enabled extended battery life and thinner, lighter mobile computers.
- Was originally intended as part of Celeron family
- 12 March 2003
- 130 nm
- 1.6 GHz
- 77 million transistors
- 1 MB integrated L2 cache

The move away from core frequency to performance begins!

Prescott

2004: Intel® Pentium® 4 Processor – 90 nm

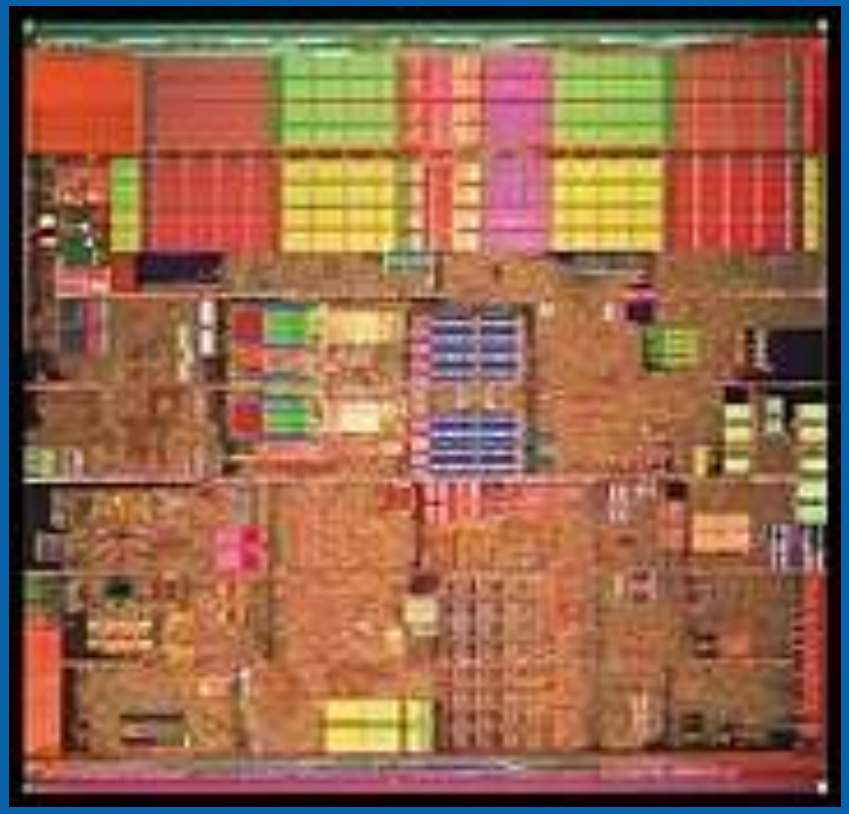

- 1MB L2 cache
- 64-bit extensions compatible with AMD64 (Humble Pie!)
- 120 million transistors
- 31 stage pipeline
- Execution Trace Cache
- 3+ GHz frequency
- ~90 Watts (Ouch!)

Frequency Push Gone Crazy!

2005: First Dual Core Opteron

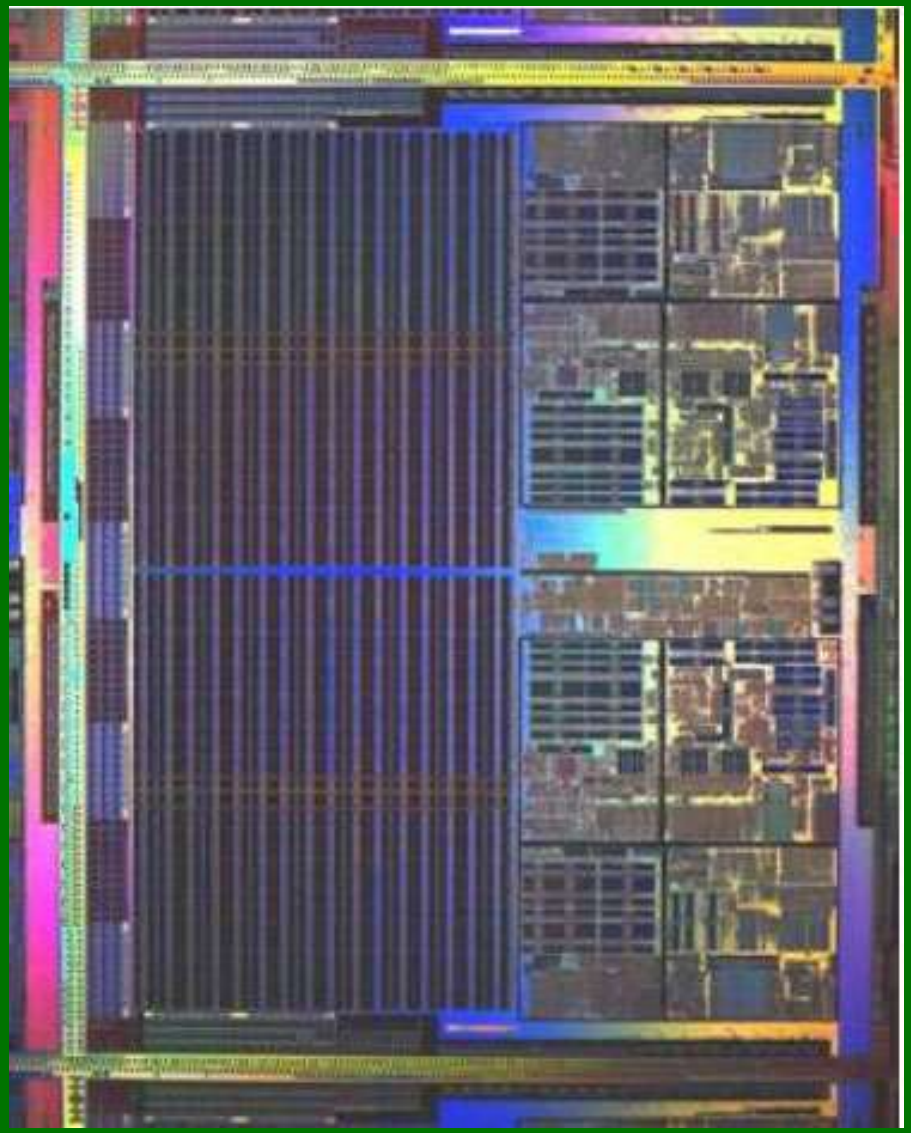

Beginning of the Multi-Core Era!

Cedar Mill

2005: Last Netburst Microarchitecture Core (65nm)

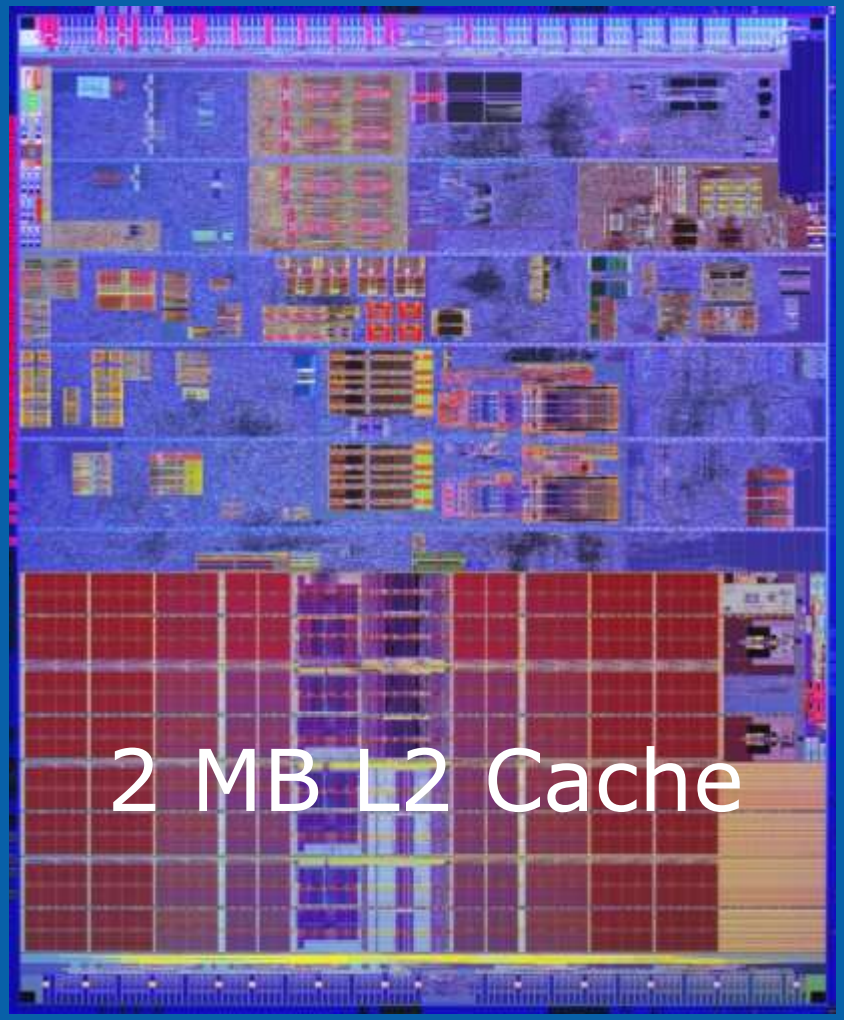

Finally the Frequency Madness Ends!

Increasing Energy Efficiency

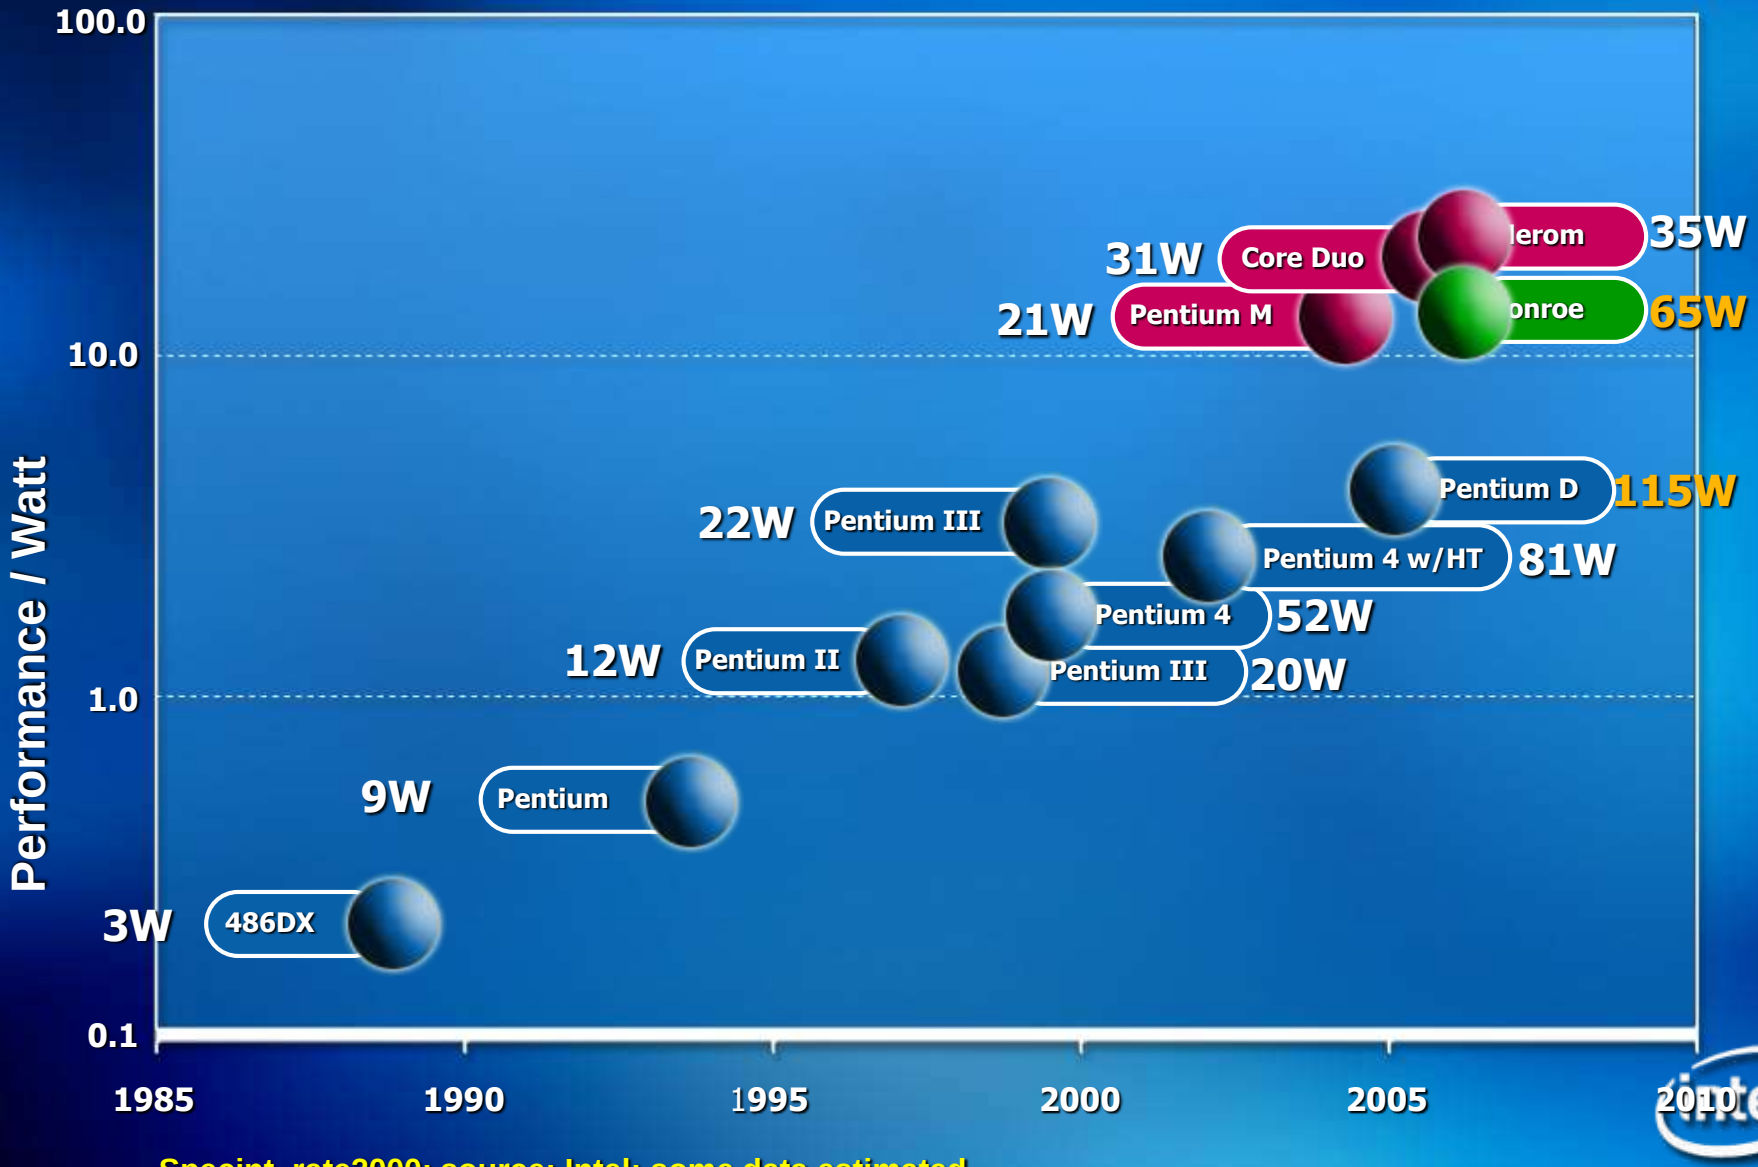

Specint_rate2000; source: Intel; some data estimated.

Yonah

2006: Intel's 1st Monolithic Dual Core

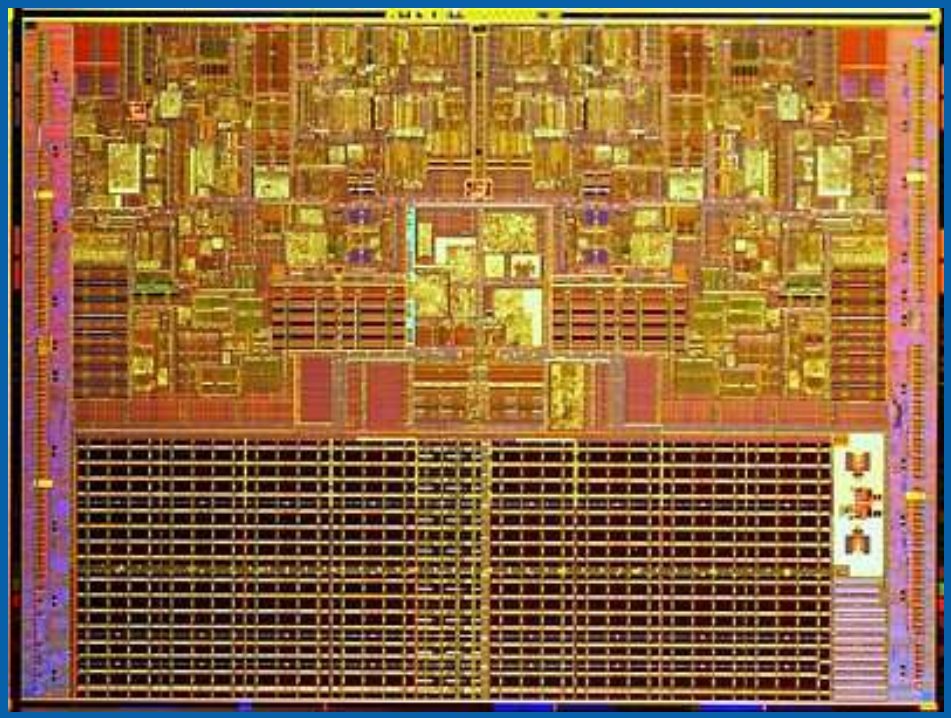

- January 2006
- Intel® Core™ Duo Processor
- 90 mm^2
- 151M transistors
- 65 nm
- First Intel processor to be used in Apple Macintosh Computers

The Convergence to Multiple Mobile Cores Begins Finally!

Why Multi-Core Processor Chips?

- With Each Process Generation transistor density doubles
  - Frequency had increased by ~1.5X; ~1.3x in future
  - Vcc had scaled by about ~0.8x; ~0.9x in future
  - Capacitance had scaled by 0.7x; ~0.8 in future
  - Total power may not scale down due to increased leakage
- Instruction Level Parallelism harder to find
- Increasing single-stream performance often requires non-linear increase in design complexity, die size, and power
- Many server applications are inherently “parallel”
- Parallelism exists in multimedia applications
- Multi-tasking usage models becoming popular

Multi-Core Energy-Efficient Performance

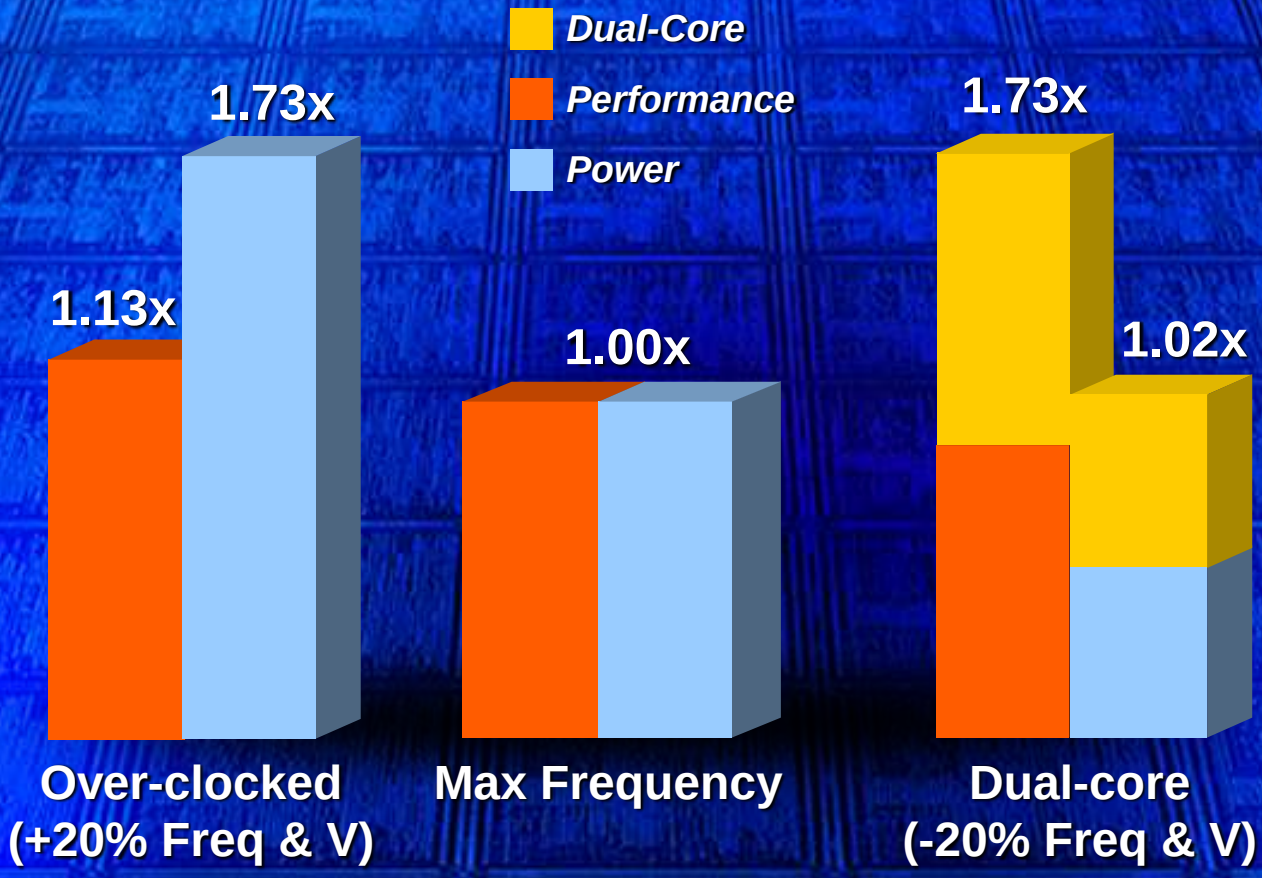

Relative single-core frequency and Vcc

2003 2004 2006 July 27, 2006

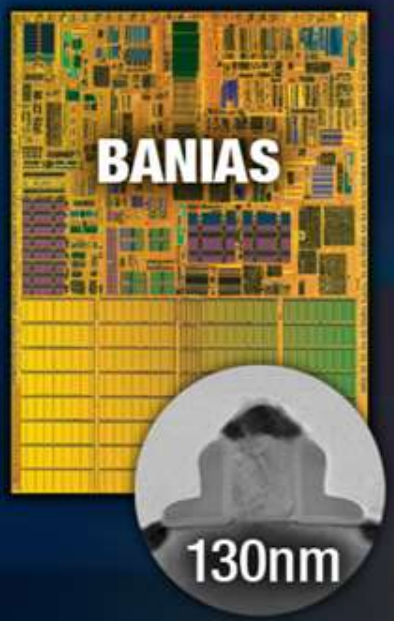

Tick

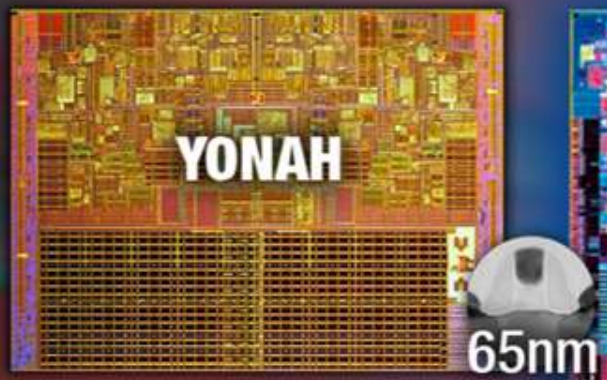

Intel® Core™ Duo Processor
90 mm²
151M transistors

Tock

Intel® Core™ 2 Duo Processor
143 mm²
291M transistors

Intel® Core™ Microarchitecture

- Intel® Wide Dynamic Execution
- Intel® Advanced Digital Media Boost
- Intel® Advanced Smart Cache
- Intel® Smart Memory Access
- Intel® Intelligent Power Capability
- Intel® 64 Architecture (Not IA-64)

intel

2006: Intel® Core™ Micro-architecture Products

- Intel® Wide Dynamic Execution
- Intel® Intelligent Power Capability
- Intel® Advanced Smart Cache
- Intel® Smart Memory Access
- Intel® Advanced Digital Media Boost

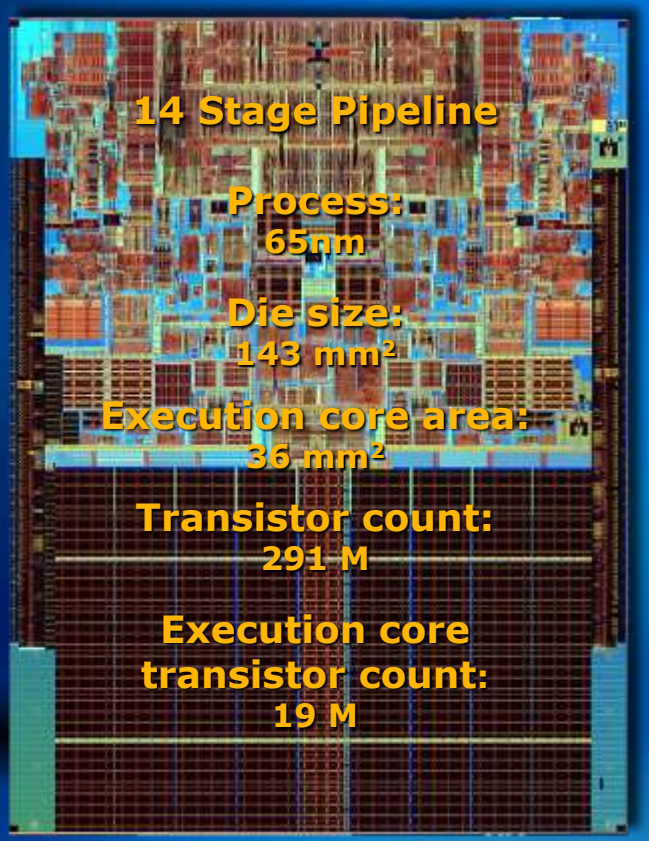

Server

Desktop

Mobile

The Empire Strikes Back!

Thanks to Israel Design Center!

Moore's Law Enables Microprocessor Advances

Chatting with Gordon Moore
http://www.youtube.com/watch?v=xzxpO0N5Amc

1.0µm 0.8µm 0.6µm 0.35µm 0.25µm 0.18µm 0.13µm 90nm 65nm

Intel 486™ Processor


Pentium® Processor




Pentium® II/III Processor




Pentium® 4 Processor



Intel® Core™ Duo Processor

Intel® Core™ 2 Duo Processor

New Designs serve High End first and waterfall to more mainstream segments as die size decreases in subsequent nodes

Source: Intel

Clovertown

October 2006: The World's First x86 Quad-Core Processor (2 die in a package)

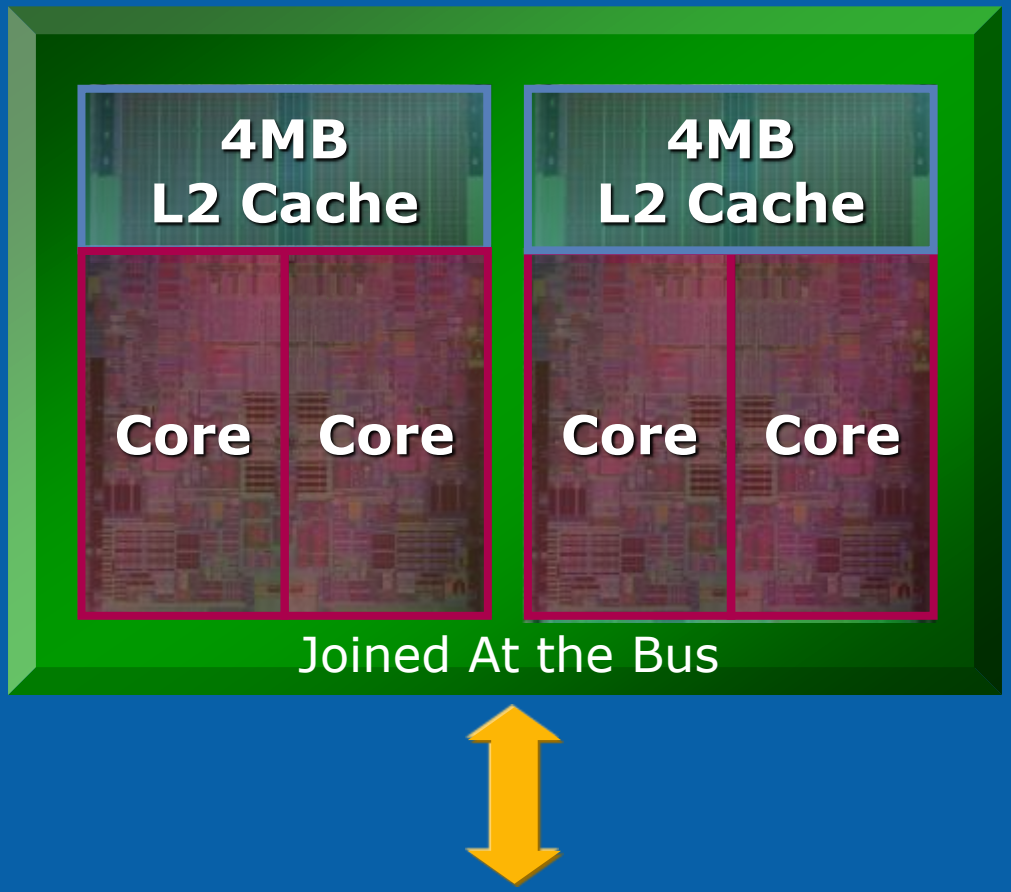

"Quick & Dirty" Innovation to drive Fast Time to Market!

2006: Itanium 2: First Billion Transistor Dual Core Chip (90nm)

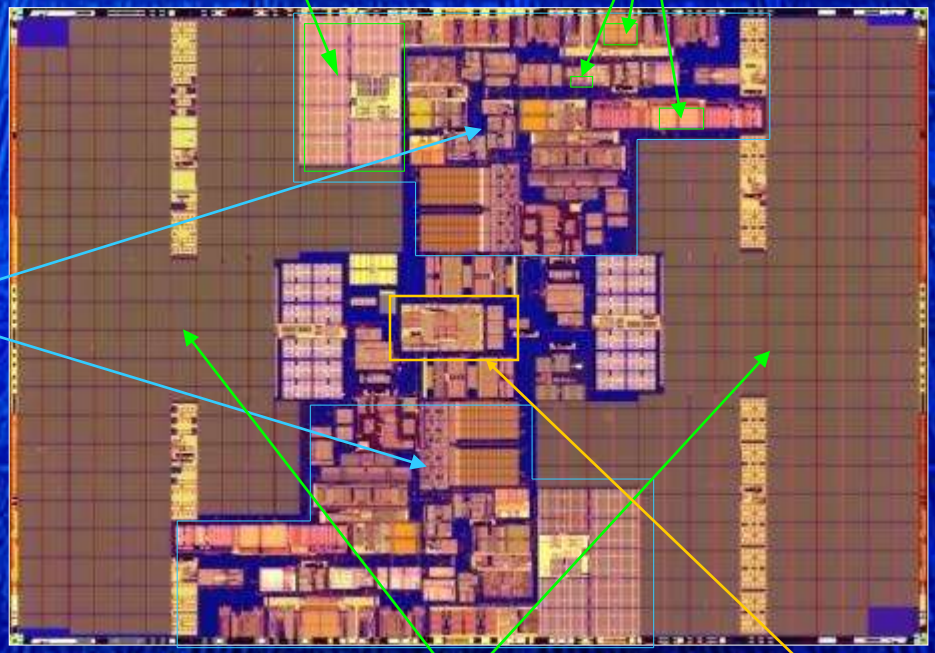

1.72 Billion Transistors (596 mm²)

From 2300 to >1Billion Transistors
In < 40 Years of Moore's Law

Moore's Law video at http://www.cs.ucr.edu/~gupta/hpca9/HPCA-PDFs/Moores_Law_Video_HPCA9.wmv

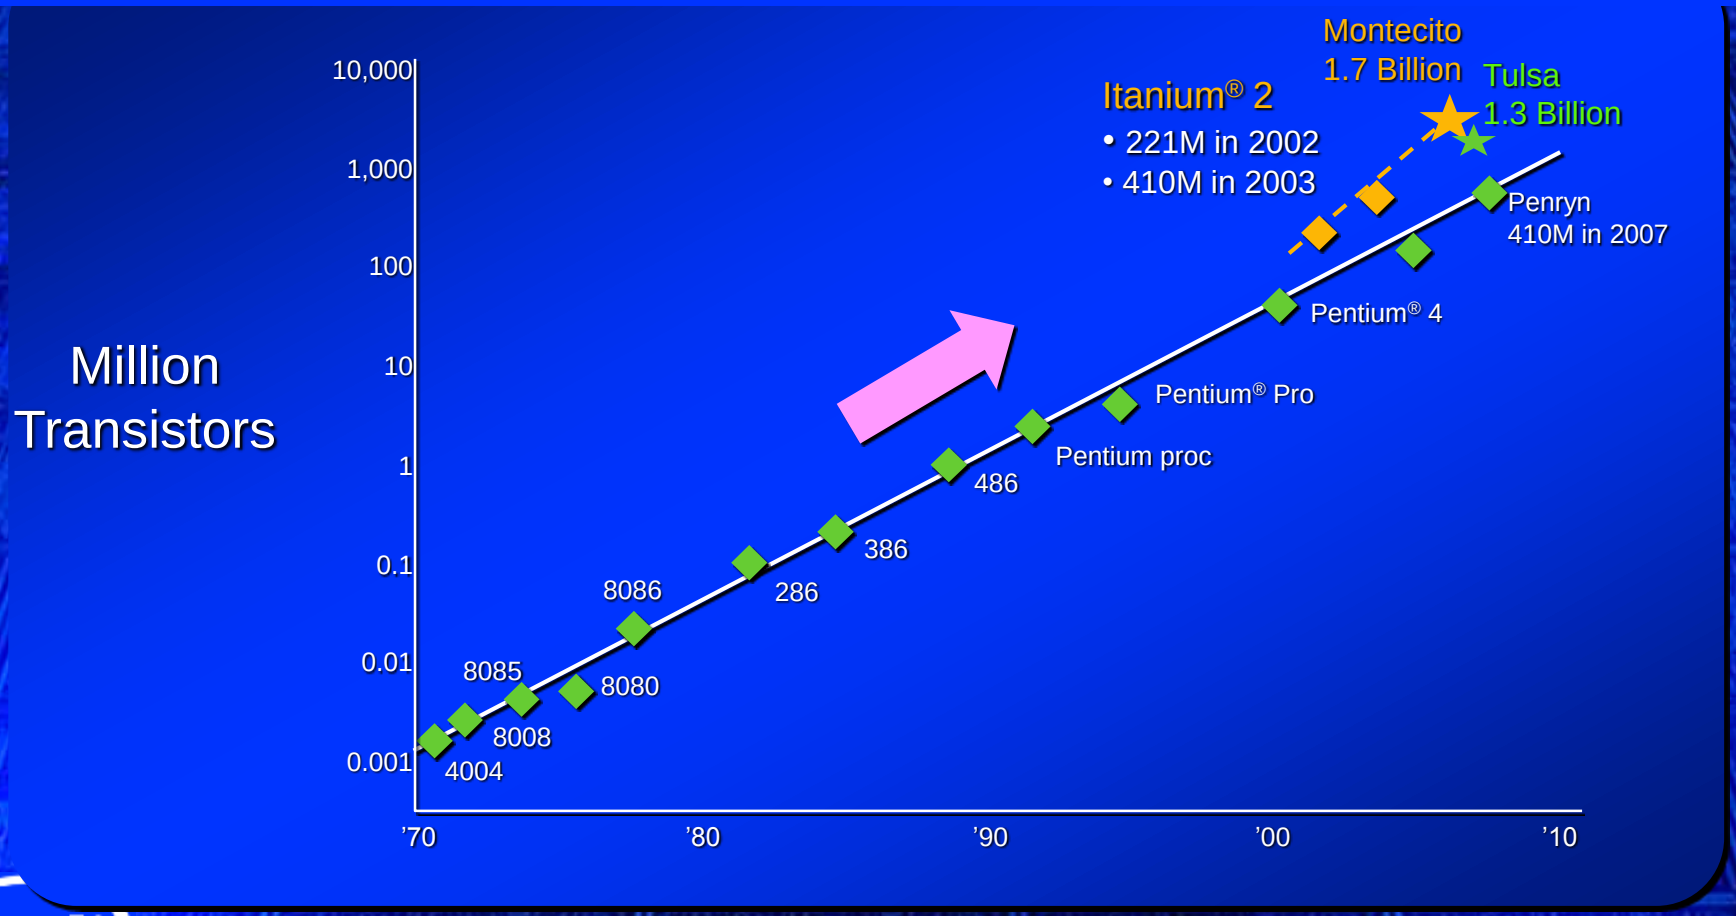

More than 1 Billion Transistors in 2006!

2007: Dual Core Penryn

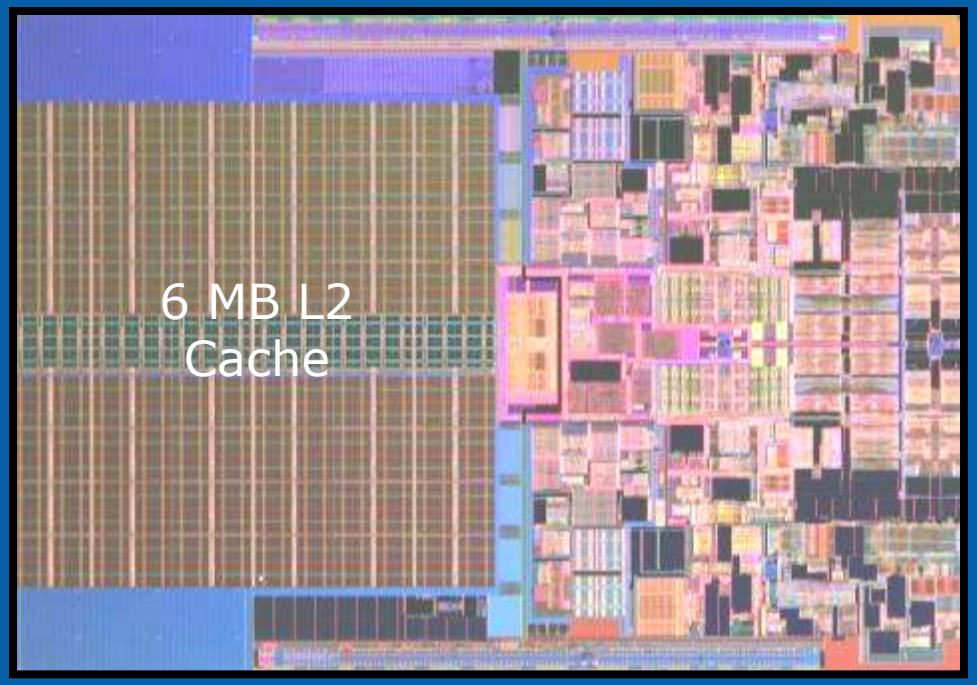

45 nm next generation Intel® Core™2 family processor

410 million transistors

World's first working 45 nm CPU

Introduced Turbo Mode

Production in the November 2007

My Last Hurrah at Intel: http://blogs.intel.com/technology/2007/04/penryn/

2007: Bill Gates Wants You!

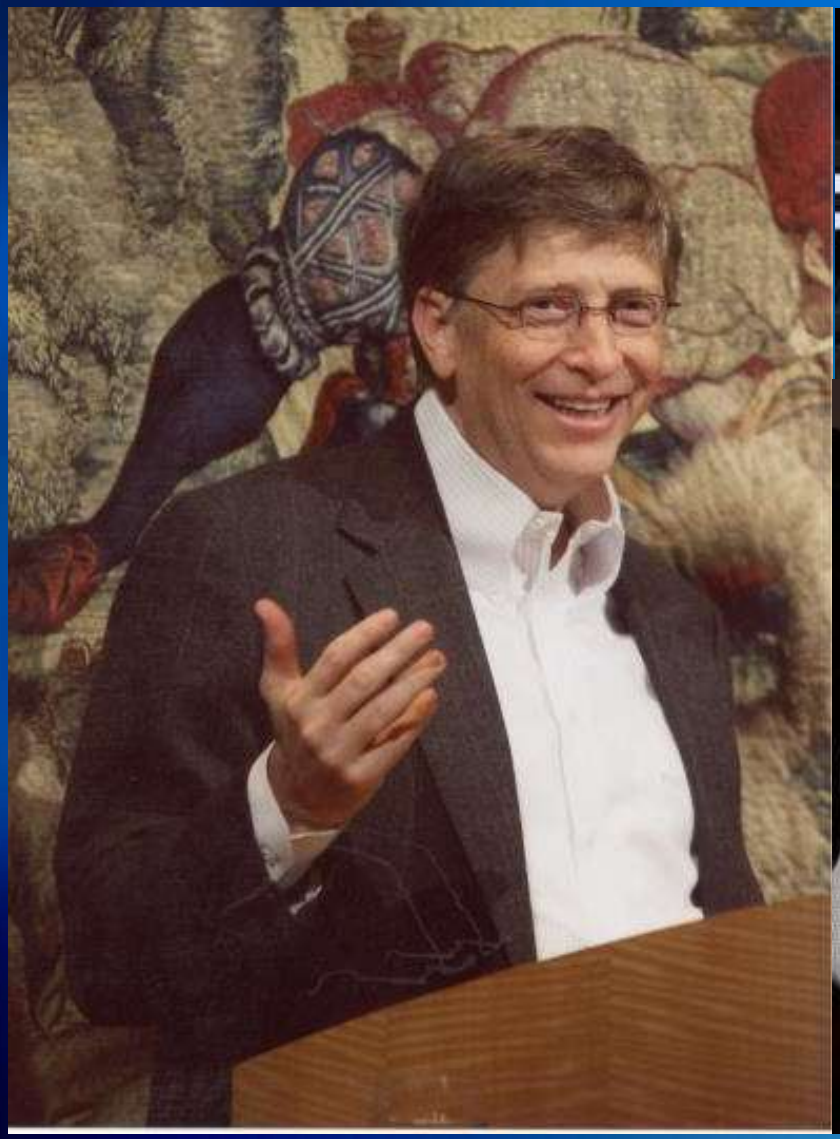

2007: AMD Barcelona
First Monolithic x86 Quad Core

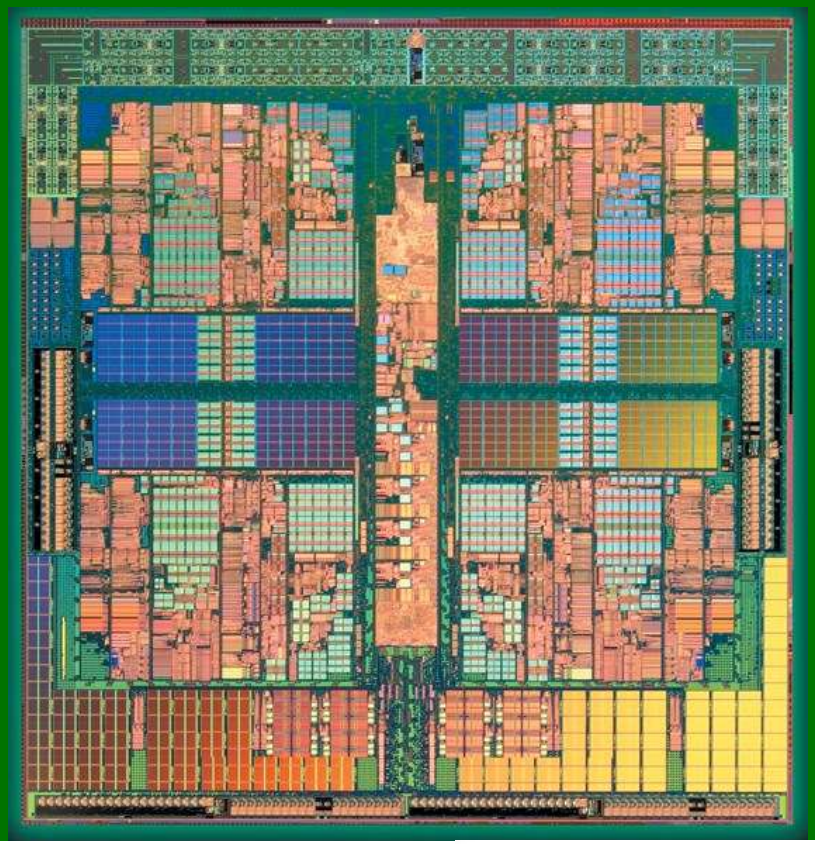

How AMD turned Barcelona into a right royal mess

Analysis The problem, the protagonists, and where to go from here

By **Charlie Demerjian**

Sat Dec 08 2007, 13:22

Chip problem limits supply of quad-core Opterons

by Scott Wasson — 1:49 PM on December 3, 2007

AMD's quad-core "Barcelona" Opterons have been notably difficult to find since their introduction two months ago, and The Tech Report has learned that a chip-level problem has impacted the supply of these chips to both server OEMs and distribution channel customers.

283mm^2 design with 463M transistors to implement four cores
and a shared 2MB L3 cache in AMD's 65nm process

2008-9: Performance Race Gets Serious With Quad Core

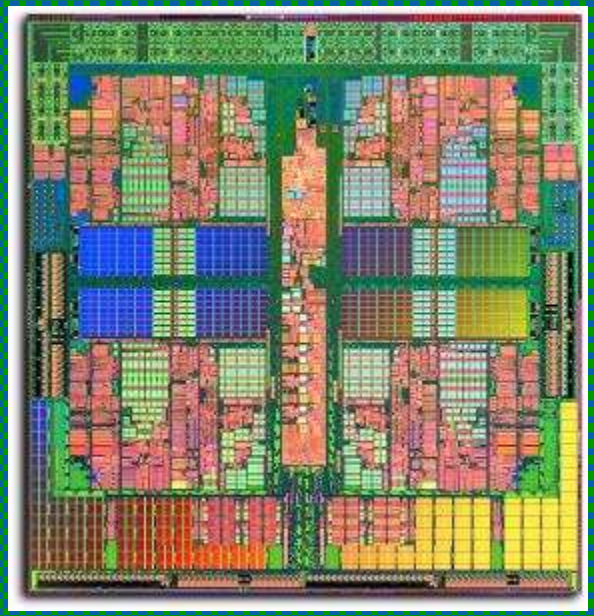

AMD Barcelona

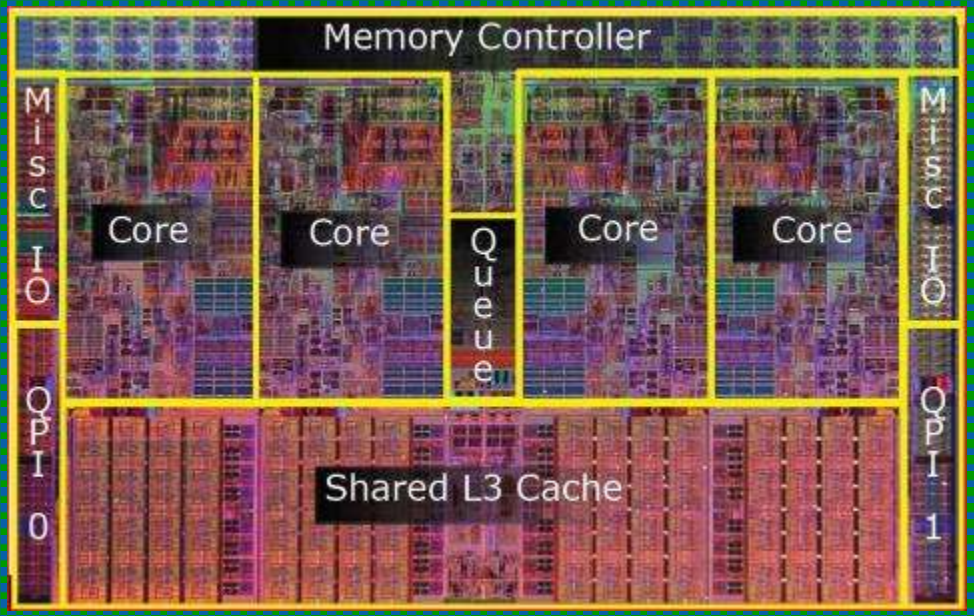

Intel Nehalem

Intel finally integrates Memory Controller and abandons shared Front Side Bus

Six Cores

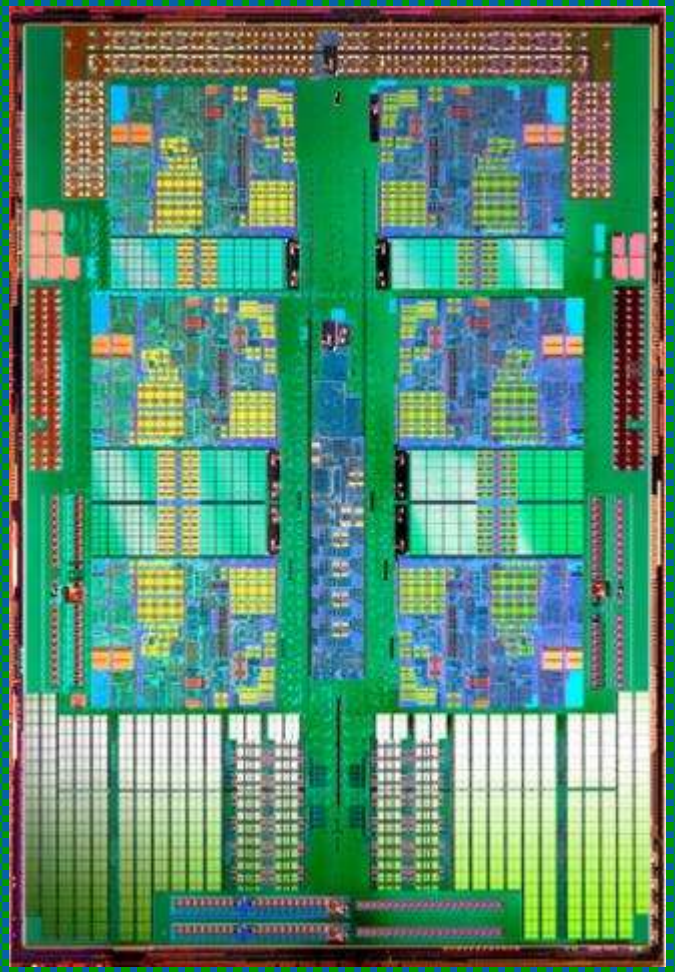

2009: AMD Istanbul

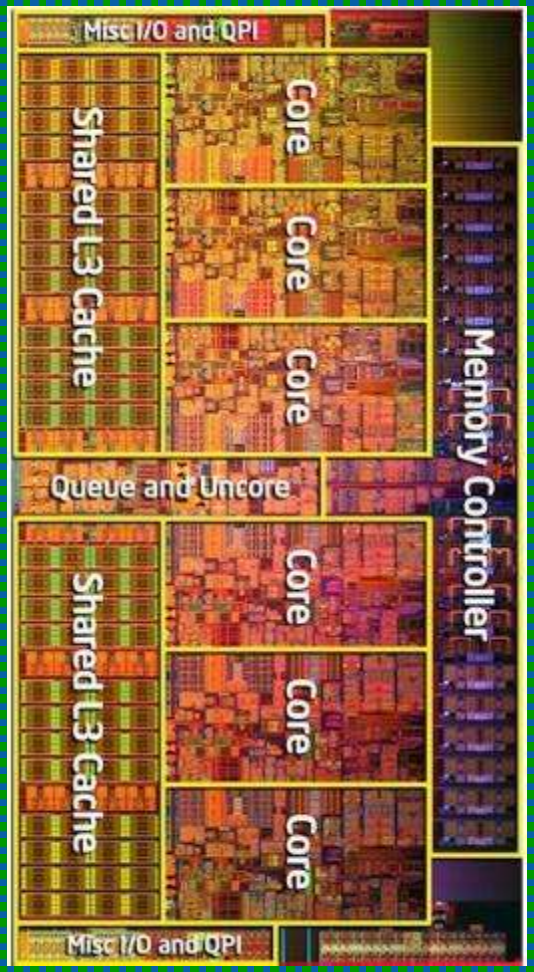

2010: Intel Westmere

Data Centers at Microsoft

The Data Center is the Computer!

Cloud Optimized High Density Servers

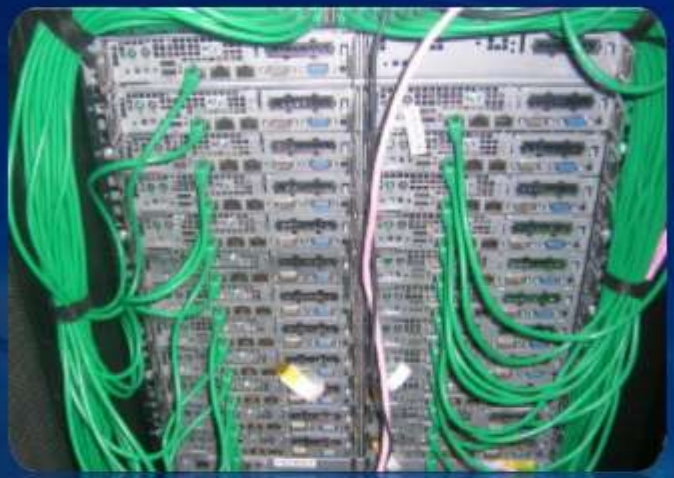

Servers built using commodity components (Low Power 2 socket CPUs, SATA HDD, MLC SSD)
No redundancy features in hardware (e.g. RAID, dual Power Supplies)
Applications specifically designed to provide Resiliency and Fault Tolerance

2013: Catching The Smartphone Wave!

QUALCOMM®

Disruptions Come from Below!

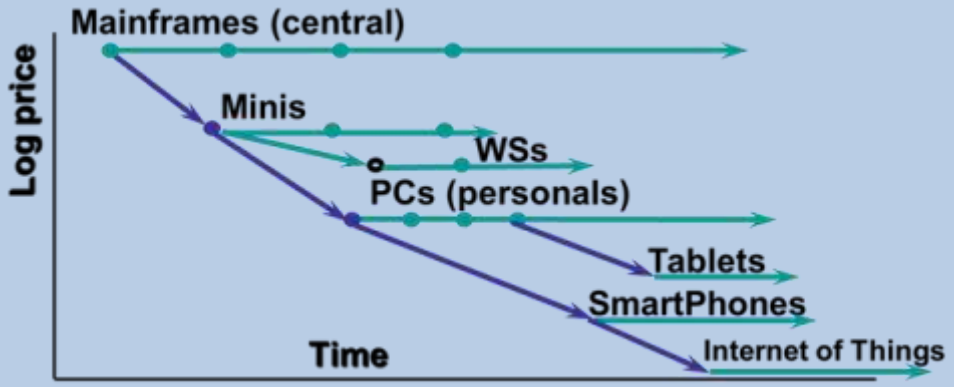

Mainframes

Minicomputers

RISC Systems

Desktop PCs

Notebooks

Smart Phones

Bell's Law:
hardware technology, networks, and interfaces allows new, smaller, more specialized computing devices to be introduced to serve a computing need.

Performance

Volume

Era of Small Cores (circa 2013)

- Intel Atom (32 nm Clover Trail)

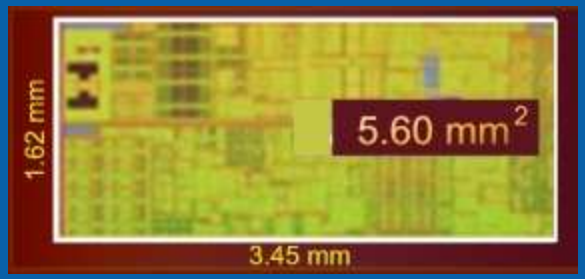

- AMD Bobcat, Jaguar (28 nm), Puma

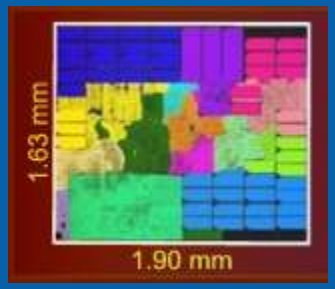

- ARM (28 nm Cortex A7 & A15)

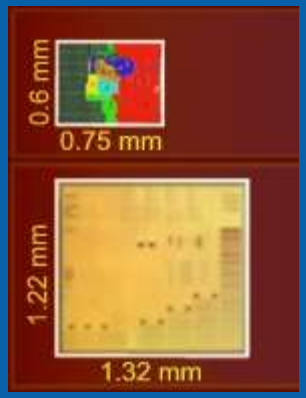

Source: www.chip-architect.com

WIRED

The Data Centers of Tomorrow Will Use the Same Tech Our Phones Do

By Peter Levine | Monday August 4, 2014

Share 0 Tweet 620 +1 0 Pin it 8

Today, the mobile phone industry is where so much innovation has been concentrated—resulting in an entirely new class of components created just for this smaller form factor: flash memory, smaller CPUs, networking hardware, and so on. Which means lightweight processors (such as ARM) and low-cost, low-power mobile components are now becoming the foundation of the next-generation datacenter.

MICROPROCESSOR report

Insightful Analysis of Processor Technology

BROADCOM BARES MUSCULAR ARM

Quad-Issue ARMv8 CPU Targets Xeon-Class Performance

By Linley Gwennap (October 21, 2013)

BROADCOM

MICROPROCESSOR report

Insightful Analysis of Processor Technology

THUNDERX RATTLES SERVER MARKET

Cavium Develops 48-Core ARM Processor to Challenge Xeon

By Linley Gwennap (June 9, 2014)

CAVIUM

Applied Micro's X-Gene challenges for server processor market

Electronics Weekly.com

Applied Micro leads the charge to infiltrate the $12 billion server processor market with ARM-based ICs.

This is not a trivial task. The $54 billion gorilla standing in Applied Micro's way is Intel with a 90% plus share of the server processor market.

So what, if any, are Applied Micro's selling points compared to Intel's?

First and foremost there's the business model.

"Competition is what we're bringing," says Gaurav Singh, vp of technical strategy at Applied Micro, "in most other markets there is very healthy competition with multiple silicon customers."

Barcelona Supercomputing Center
Centro Nacional de Supercomputación

EUROPE WANTS A SMARTPHONE SUPERCOMPUTER

A consortium hopes to build exaflop supercomputers from mobile CPUs

Intel juices up microserver speeds with thrifty Avoton chip

intel Atom

Summary: *Intel is claiming to have made significant strides in performance and power efficiency in the microserver market with its new Avoton system on a chip.*

ZDNet

AMD Announces the Availability of 64-bit ARM Opteron Developer Kits

SUNNYVALE, Calif. 7/30/2014

AMD (NYSE: AMD) today announced the immediate availability of the AMD Opteron™ A1100-Series developer kit, which features AMD's first 64-bit ARM®-based processor, codenamed "Seattle." AMD is the first company to provide a standard ARM Cortex®-A57- based server platform for software developers and integrators. Software and hardware developers as well as early adopters in large datacenters are eligible and can apply on AMD's website.

"The journey toward a more efficient infrastructure for large-scale datacenters is taking a major step forward today with broader availability of our AMD Opteron A1100-Series development kit," said Suresh Gopalakrishnan, general manager and vice president, server business unit, AMD. "After successfully sampling to major ecosystem partners such as firmware, OS, and tools providers, we are taking the next step in what will be a collaborative effort across the industry to reimagine the datacenter based on the open business model of ARM innovation."

With this announcement, AMD becomes the only provider of 64-bit ARM server hardware with complete ARMv8 instruction support to foster the development of the ecosystem for efficient storage, Web applications and hosting. AMD is the only provider to offer the standard ARM Cortex-A57 technology.

Contact:

Kristen Lisa
AMD Public Relations
(512) 602-6633
kristen.lisa@amd.com

AMD

The Smart Phone Era Is Redefining Computing

"The phone in your pocket will be as much of a computer as anyone needs".
– Dr. Irwin Jacobs, 2000

PC Market Shift

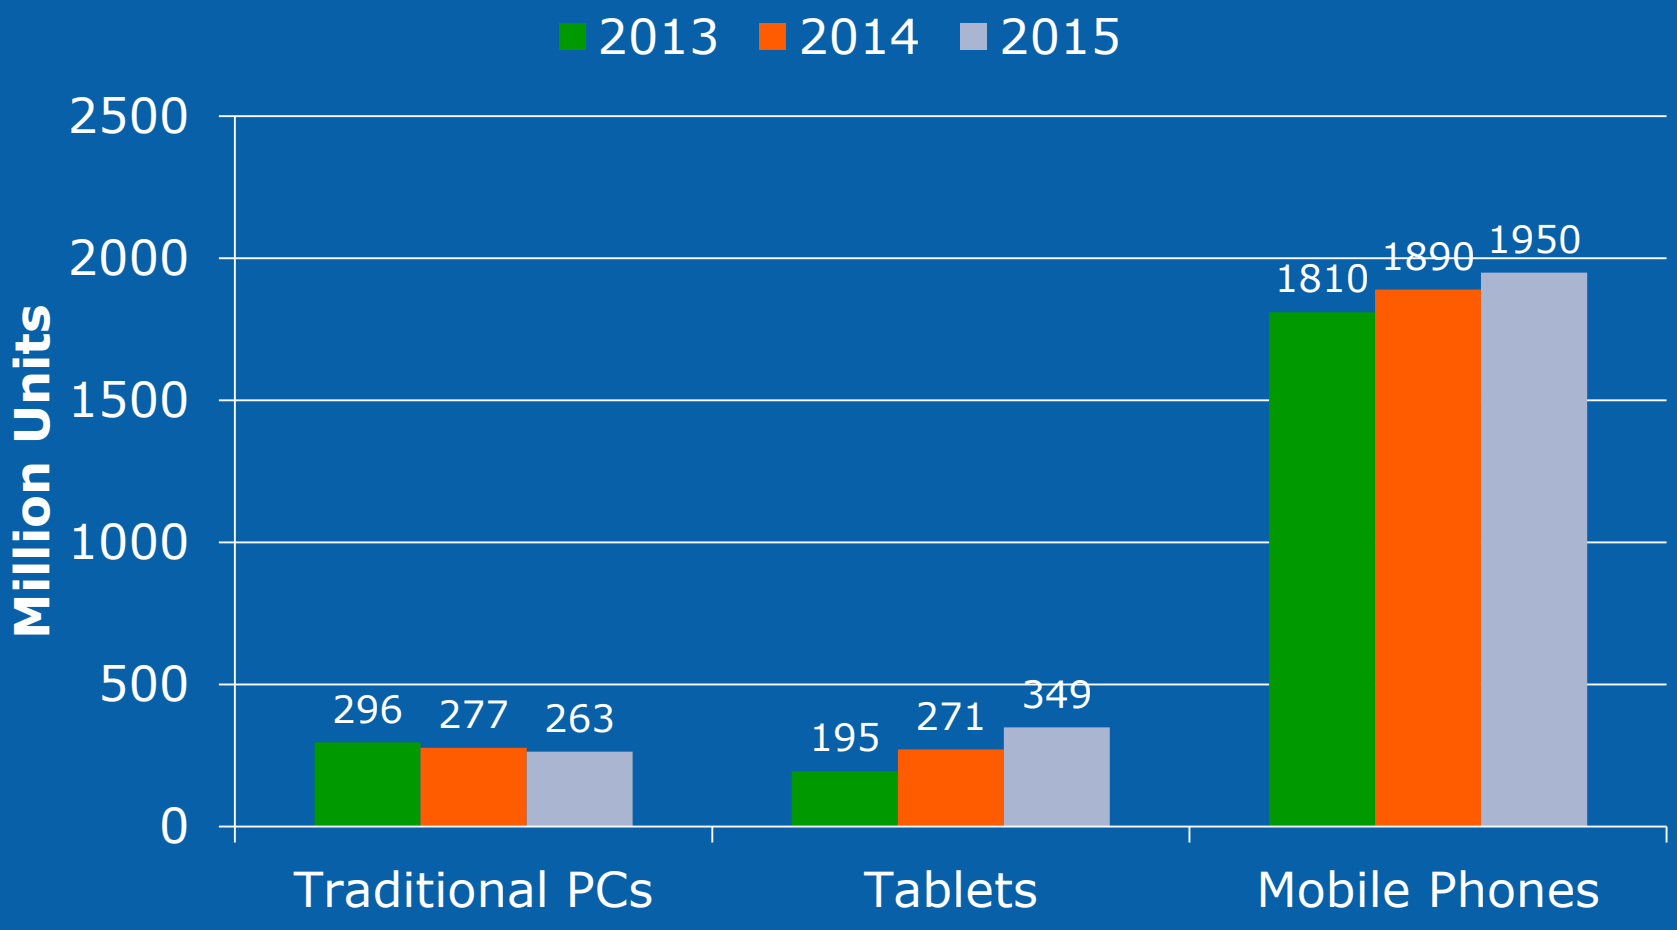

Source: www.pctoday.com

Continued Smartphone Momentum

~20% CAGR for smartphone unit shipments expected between 2012-2017

~7B

Cumulative smartphone unit shipments forecast between 2013-2017

Source: Gartner, Sept '13

Smartphone System Architecture

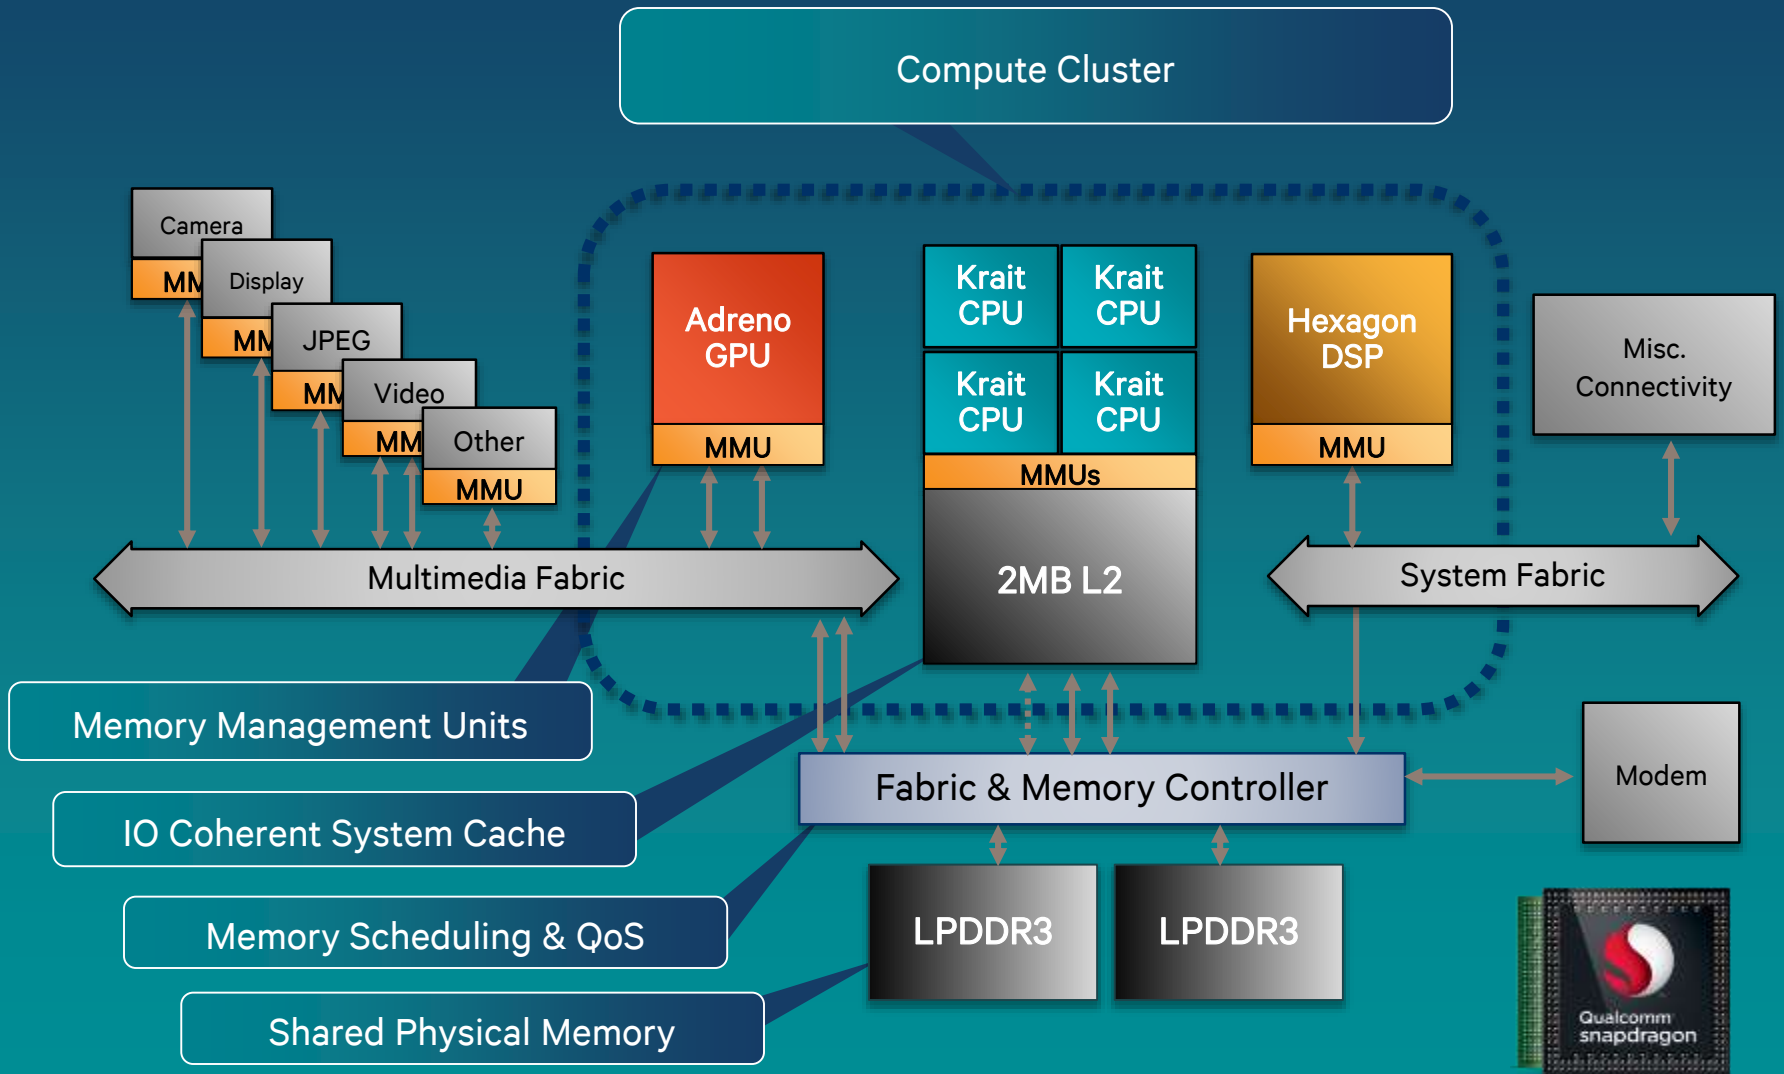

Technology Cycles – Still Early Cycle on Smartphones + Tablets, Now Wearables Coming on Strong, Faster than Typical 10-Year Cycle

Technology Cycles Have Tended to Last Ten Years

Mainframe Computing 1960s	*Mini* Computing 1970s	*Personal* Computing 1980s	*Desktop Internet* Computing 1990s	*Mobile Internet* Computing 2000s	*Wearable / Everywhere* Computing 2014+
				 	   Others?

Image Source: Computersciencelab.com, Wikipedia, IBM, Apple, Google, NTT docomo, Google, Jawbone, Pebble.

Learn to Wear Many Hats!

“Don’t be encumbered by past history, go off and do something wonderful.”
- Bob Noyce, Intel Founder


Questions?