Parallelism in Mainstream Enterprise Platforms of the Future

Dileep Bhandarkar
Architect at Large
Enterprise Platforms Group
Intel Corporation

September 23rd, 2002

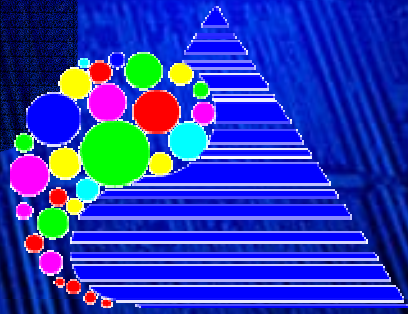

Outline

- Semiconductor Technology Evolution
- Moore's Law Video
- Parallelism in Microprocessors Today
- Multiprocessor Systems
- The Billion Transistor Chip
- Summary

Birth of the Revolution -- The Intel 4004

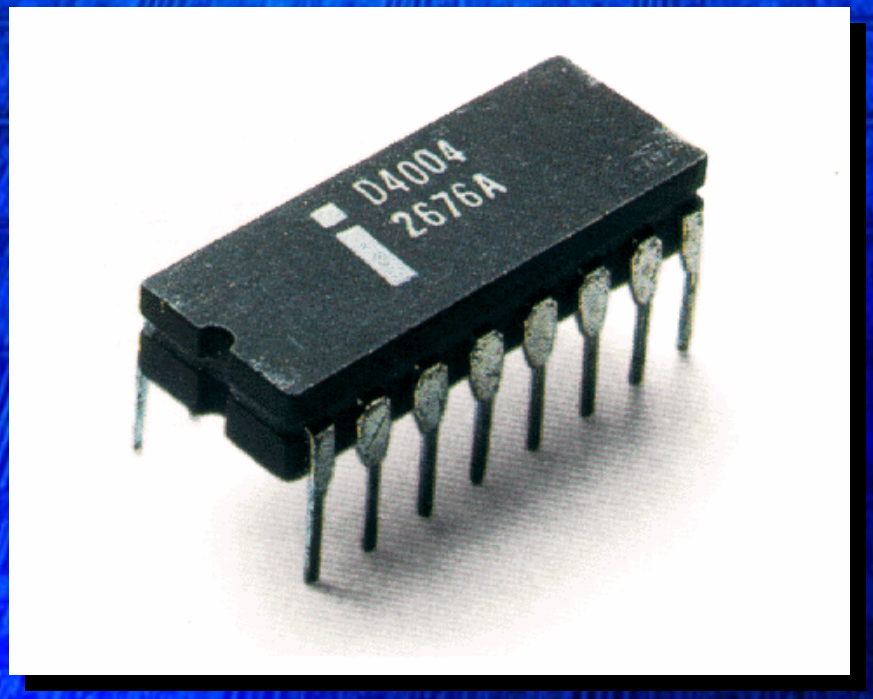

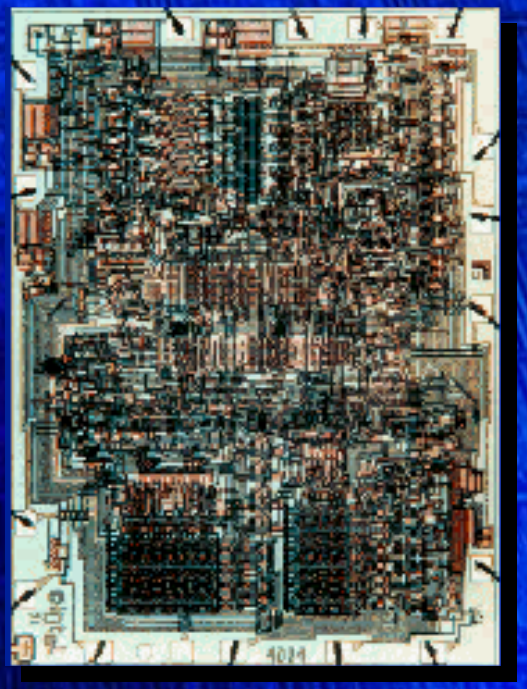

Introduced November 15, 1971
108 KHz, 50 KIPs , 2300 10μ transistors

2001 – Pentium® 4 Processor

Introduced November 20, 2000
@1.5 GHz core, 400 MT/s bus

August 27, 2001
@2 GHz, 400 MT/s bus
640 SPECint_base2000
704 SPECfp_base2000

42 Million 0.18µ transistors

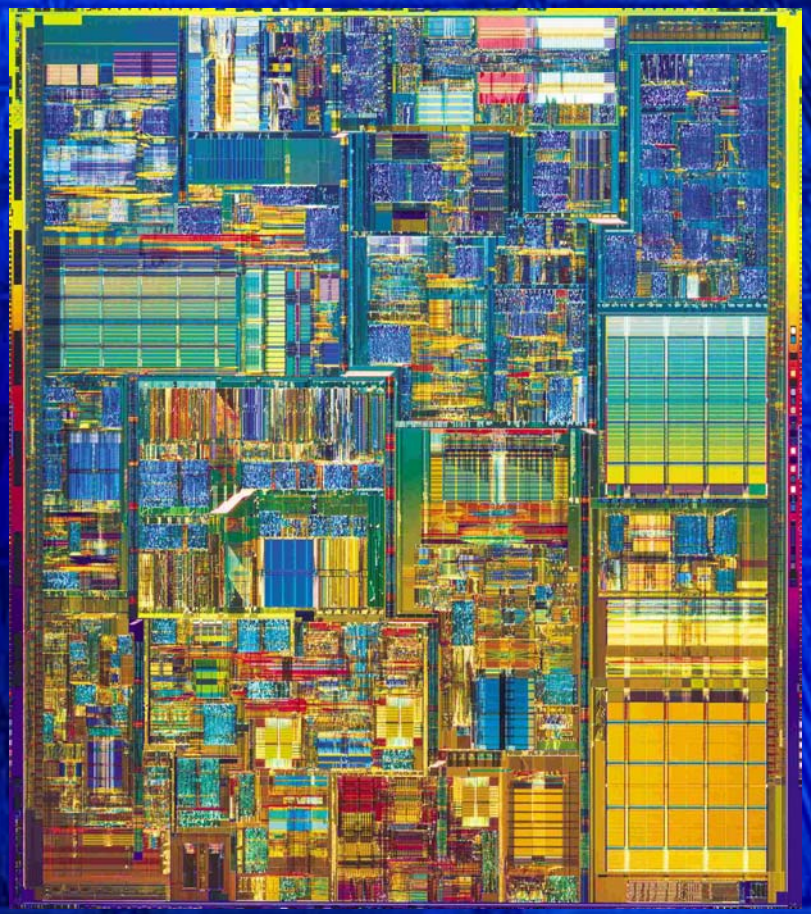

30 Years of Progress

- 4004 to Pentium® 4 processor
  - Transistor count: 20,000x increase
  - Frequency: 20,000x increase
  - 39% Compound Annual Growth rate

2002 – Pentium® 4 Processor

August 26, 2002
@2.8 GHz, 533 MT/s bus
976 SPECint_base2000
915 SPECfp_base2000
55 Million 130 nm process

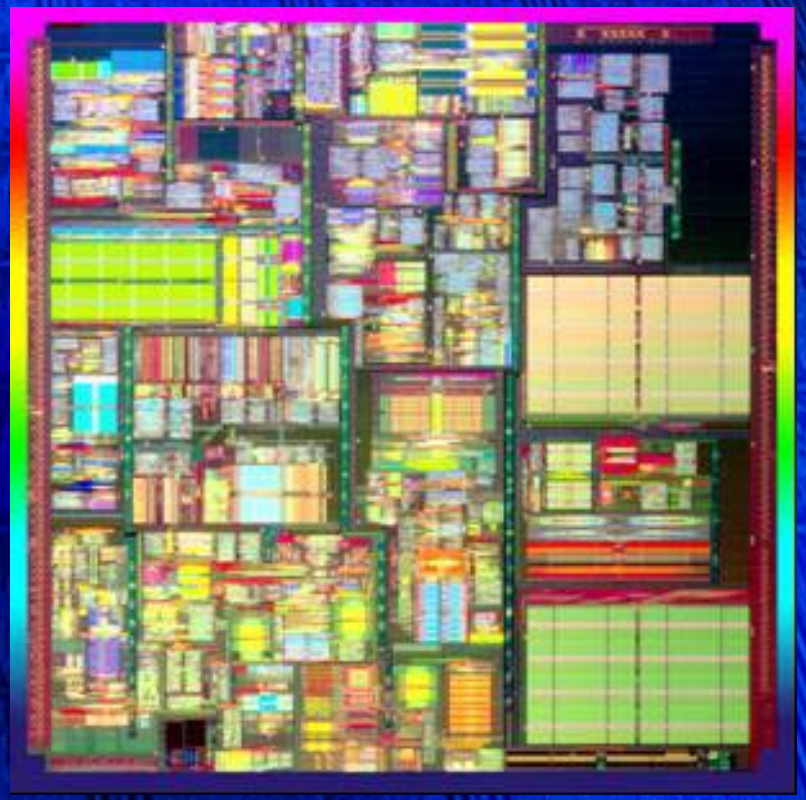

intel

Itanium® 2 Processor Overview

- .18µm bulk, 6 layer Al process
- 8 stage, fully stalled in-order pipeline
- Symmetric six integer-issue design
- IA32 execution engine integrated
- 3 levels of cache on-die totaling 3.3MB
- 221 Million transistors
- 130W @1GHz, 1.5V

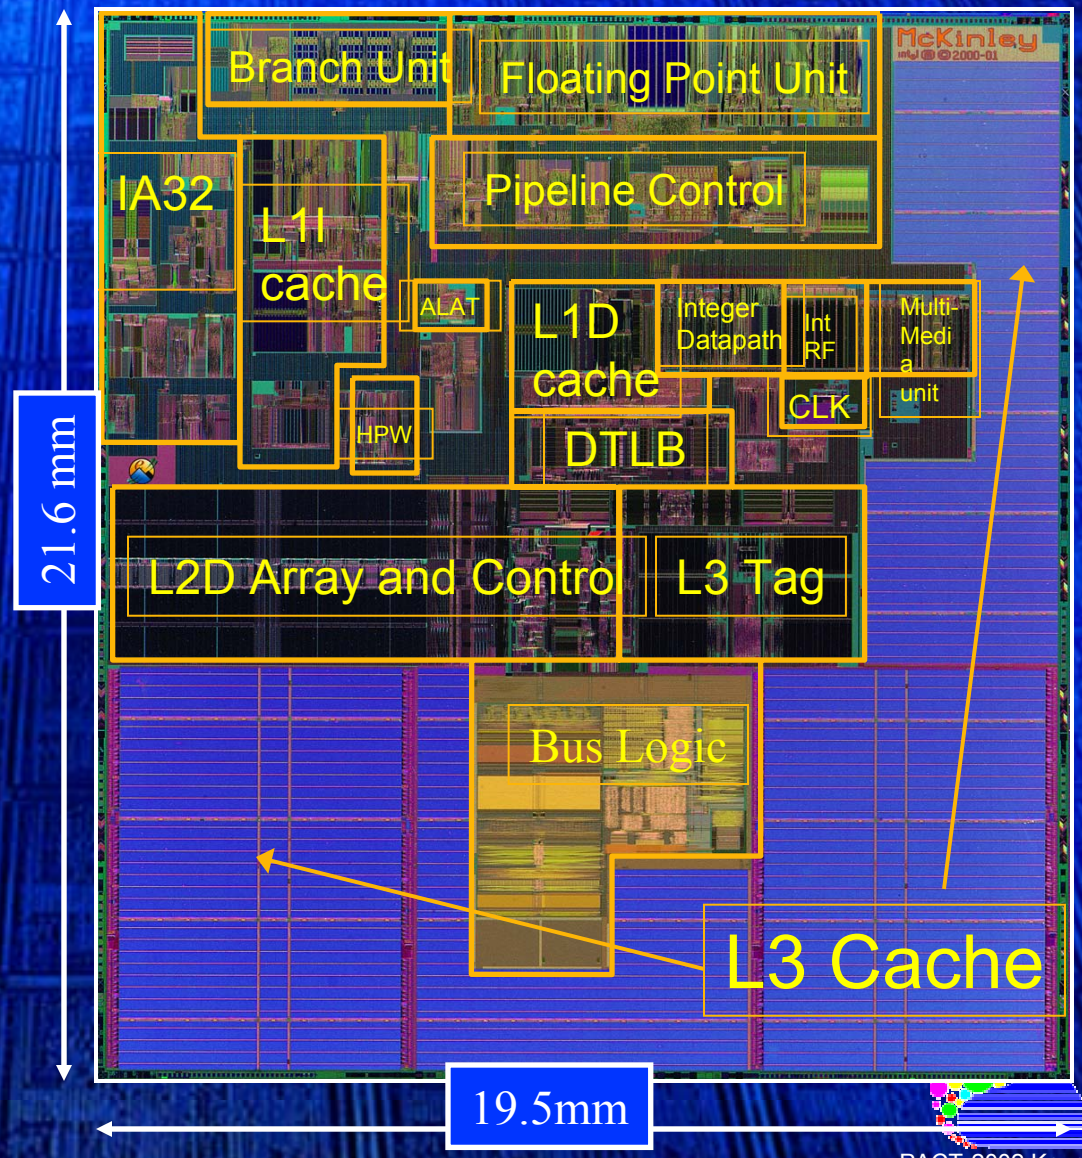

Continuing at this Rate by End of the Decade

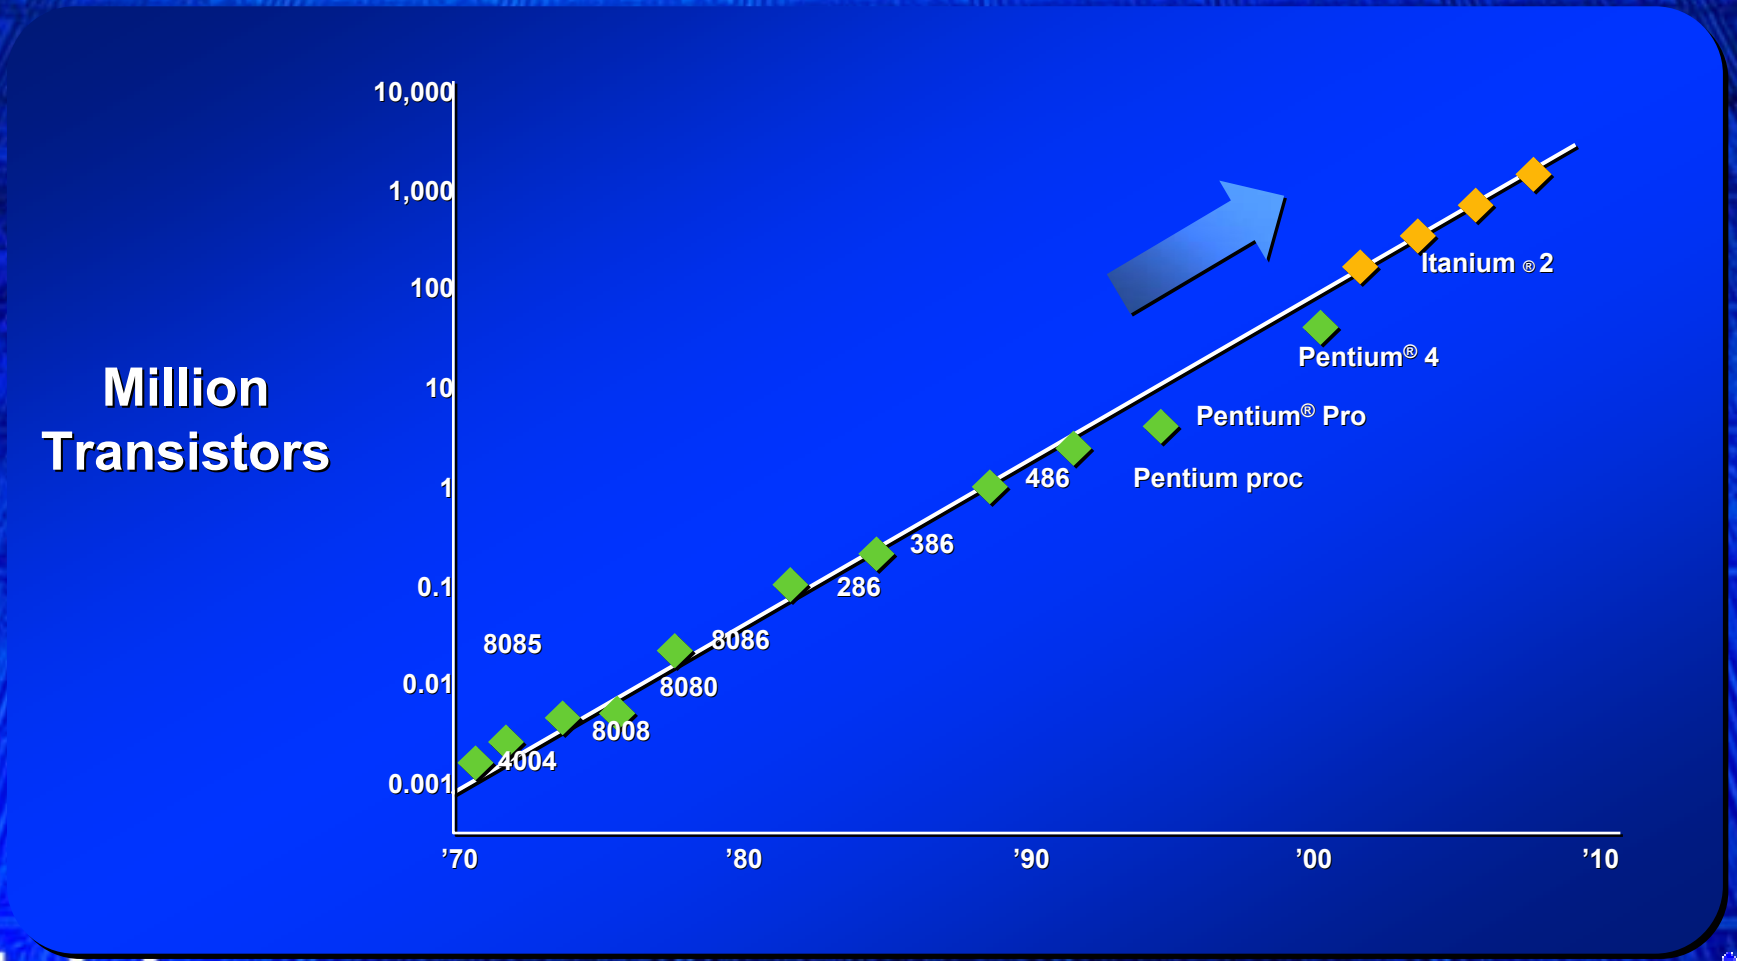

Billion Transistors by 2005

"If the automobile industry advanced as rapidly as the semiconductor industry, a Rolls Royce would get 1/2 million miles per gallon and it would be cheaper to throw it away than to park it."

Gordon Moore,
Intel Corporation

Semiconductor Manufacturing Process Evolution

	Actual					Forecast	
Process name	P852	P854	P856	P858	Px60	P1262	P1264
Production	1993	1995	1997	1999	2001	2003	2005
Generation	0.50	0.35	0.25	0.18µm	130 nm	90 nm	65 nm
Gate Length	0.50	0.35	0.20	0.13	<70 nm	<50 nm	<35 nm
Wafer Size (mm)	200	200	200	200	200/300	300	300

New generation every 2 years

Outline

- Semiconductor Technology Evolution
- Moore's Law Video
- Parallelism in Microprocessors Today
- Multiprocessor Systems
- The Billion Transistor Chip
- Summary

Moore's Law

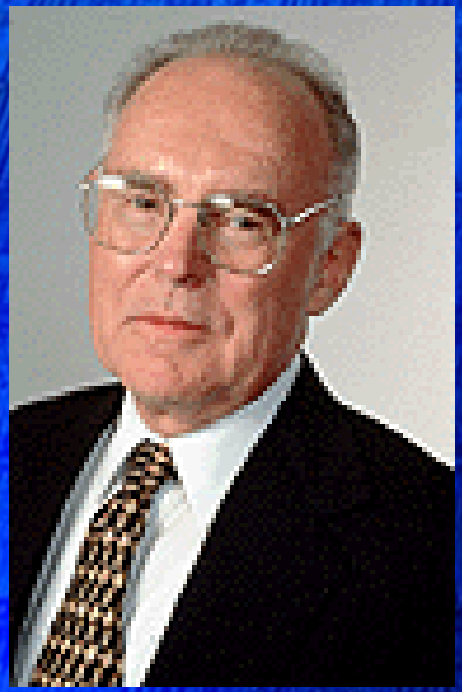

The experts look ahead

Cramming more components onto integrated circuits

With unit cost falling as the number of components per circuit rises, by 1975 economics may dictate squeezing as many as 65,000 components on a single silicon chip

By Gordon E. Moore

Director, Research and Development Laboratories, Fairchild Semiconductor division of Fairchild Camera and Instrument Corp.

The future of integrated electronics is the future of electronics itself. The advantages of integration will bring about a proliferation of electronics, pushing this science into many new areas.

Integrated circuits will lead to such wonders as home computers—or at least terminals connected to a central computer—automatic controls for automobiles, and personal portable communications equipment. The electronic wristwatch needs only a display to be feasible today.

But the biggest potential lies in the production of large systems. In telephone communications, integrated circuits in digital filters will separate channels on multiplex equipment. Integrated circuits will also switch telephone circuits and perform data processing.

Computers will be more powerful, and will be organized in completely different ways. For example, memories built of integrated electronics may be distributed throughout the machine instead of being concentrated in a central unit. In addition, the improved reliability made possible by integrated circuits will allow the construction of larger processing units. Machines similar to those in existence today will be built at lower costs and with faster turn-around.

Present and future

By integrated electronics, I mean all the various technologies which are referred to as microelectronics today as well as any additional ones that result in electronics functions supplied to the user as irreducible units. These technologies were first investigated in the late 1950's. The object was to miniaturize electronics equipment to include increasingly complex electronic functions in limited space with minimum weight. Several approaches evolved, including microassembly techniques for individual components, thin-film structures and semiconductor integrated circuits.

Each approach evolved rapidly and converged so that each borrowed techniques from another. Many researchers believe the way of the future to be a combination of the various approaches.

The advocates of semiconductor integrated circuitry are already using the improved characteristics of thin-film resistors by applying such films directly to the active semiconductor substrate. Those advocating a technology based upon films are developing sophisticated techniques for the attachment of active semiconductor devices to the passive film arrays.

Both approaches have worked well and are being used in equipment today.

The author

Dr. Gordon E. Moore is one of the new breed of electronic engineers, schooled in the physical sciences rather than in electronics. He earned a B.S. degree in chemistry from the University of California and a Ph.D. degree in physical chemistry from the California Institute of Technology. He was one of the founders of Fairchild Semiconductor and has been director of the research and development laboratories since 1959.

Electronics, Volume 38, Number 8, April 19, 1965

Outline

- Semiconductor Technology Evolution
- Moore's Law Video
- Parallelism in Microprocessors Today
- Multiprocessor Systems
- The Billion Transistor Chip
- Summary

Parallelism at Multiple Levels

- Within a processor
  - multiple issue processors with lots of execution units
  - wider superscalar
  - explicit parallelism
- Multiple processors on a chip
  - Hardware Multi Threading
  - Multiple cores
- System Level Multiprocessors

EPIC Architecture Features

Explicitly Parallel Instruction Computing

- Enable wide execution by providing processor implementations that compiler can take advantage of
- Performance through parallelism
  - Multiple execution units and issue ports in parallel
  - 2 bundles (up to 6 Instructions) dispatched every cycle
- Massive on-chip resources
  - 128 general registers, 128 floating point registers
  - 64 predicate registers, 8 branch registers
  - Exploit parallelism
  - Efficient management engines (register stack engine)
- Provide features that enable compiler to reschedule programs using advanced features (predication, speculation)
- Enable, enhance, express, and exploit parallelism

Instruction Formats: Bundles

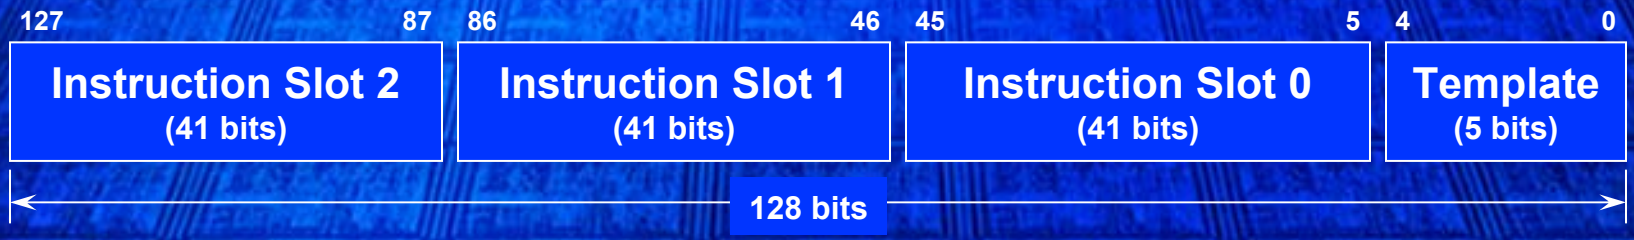

- Template identifies types of instructions in bundle and delineates independent operations (through "stops")
- Instruction types
  - M: Memory
  - I: Shifts and multimedia
  - A: ALU
  - B: Branch
  - F: Floating point
  - L+X: Long

- Template encodes types
  - MII, MLX, MMI, MFI, MMF, MI_I, M_MI
  - Branch: MIB, MMB, MFB, MBB, BBB
- Template encodes parallelism
  - All come in two flavors: with and without stop at end

Itanium® 2 Processor Architecture

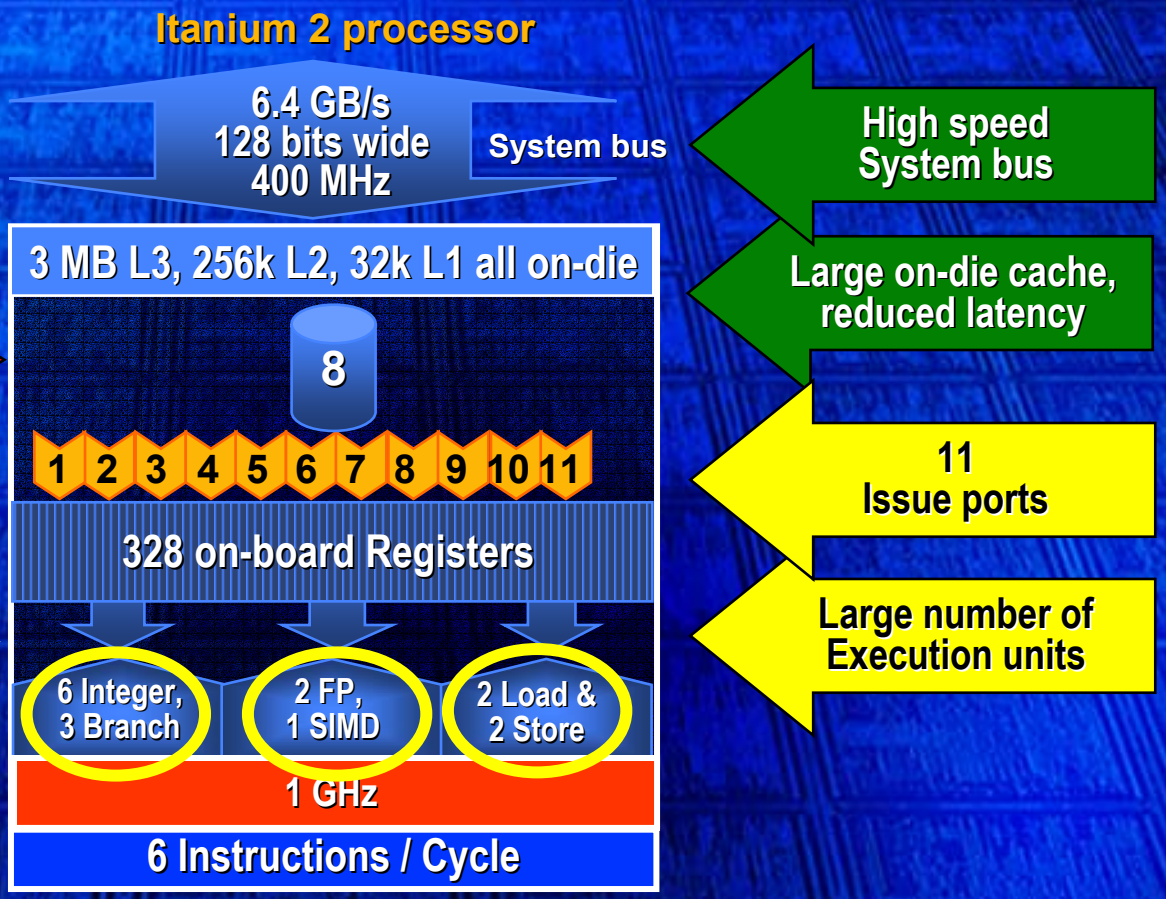

Lots of parallelism within a single core

Itanium 2 processor
221 million transistors total
25 million in CPU core

Processor Structure

McKinley Block Diagram

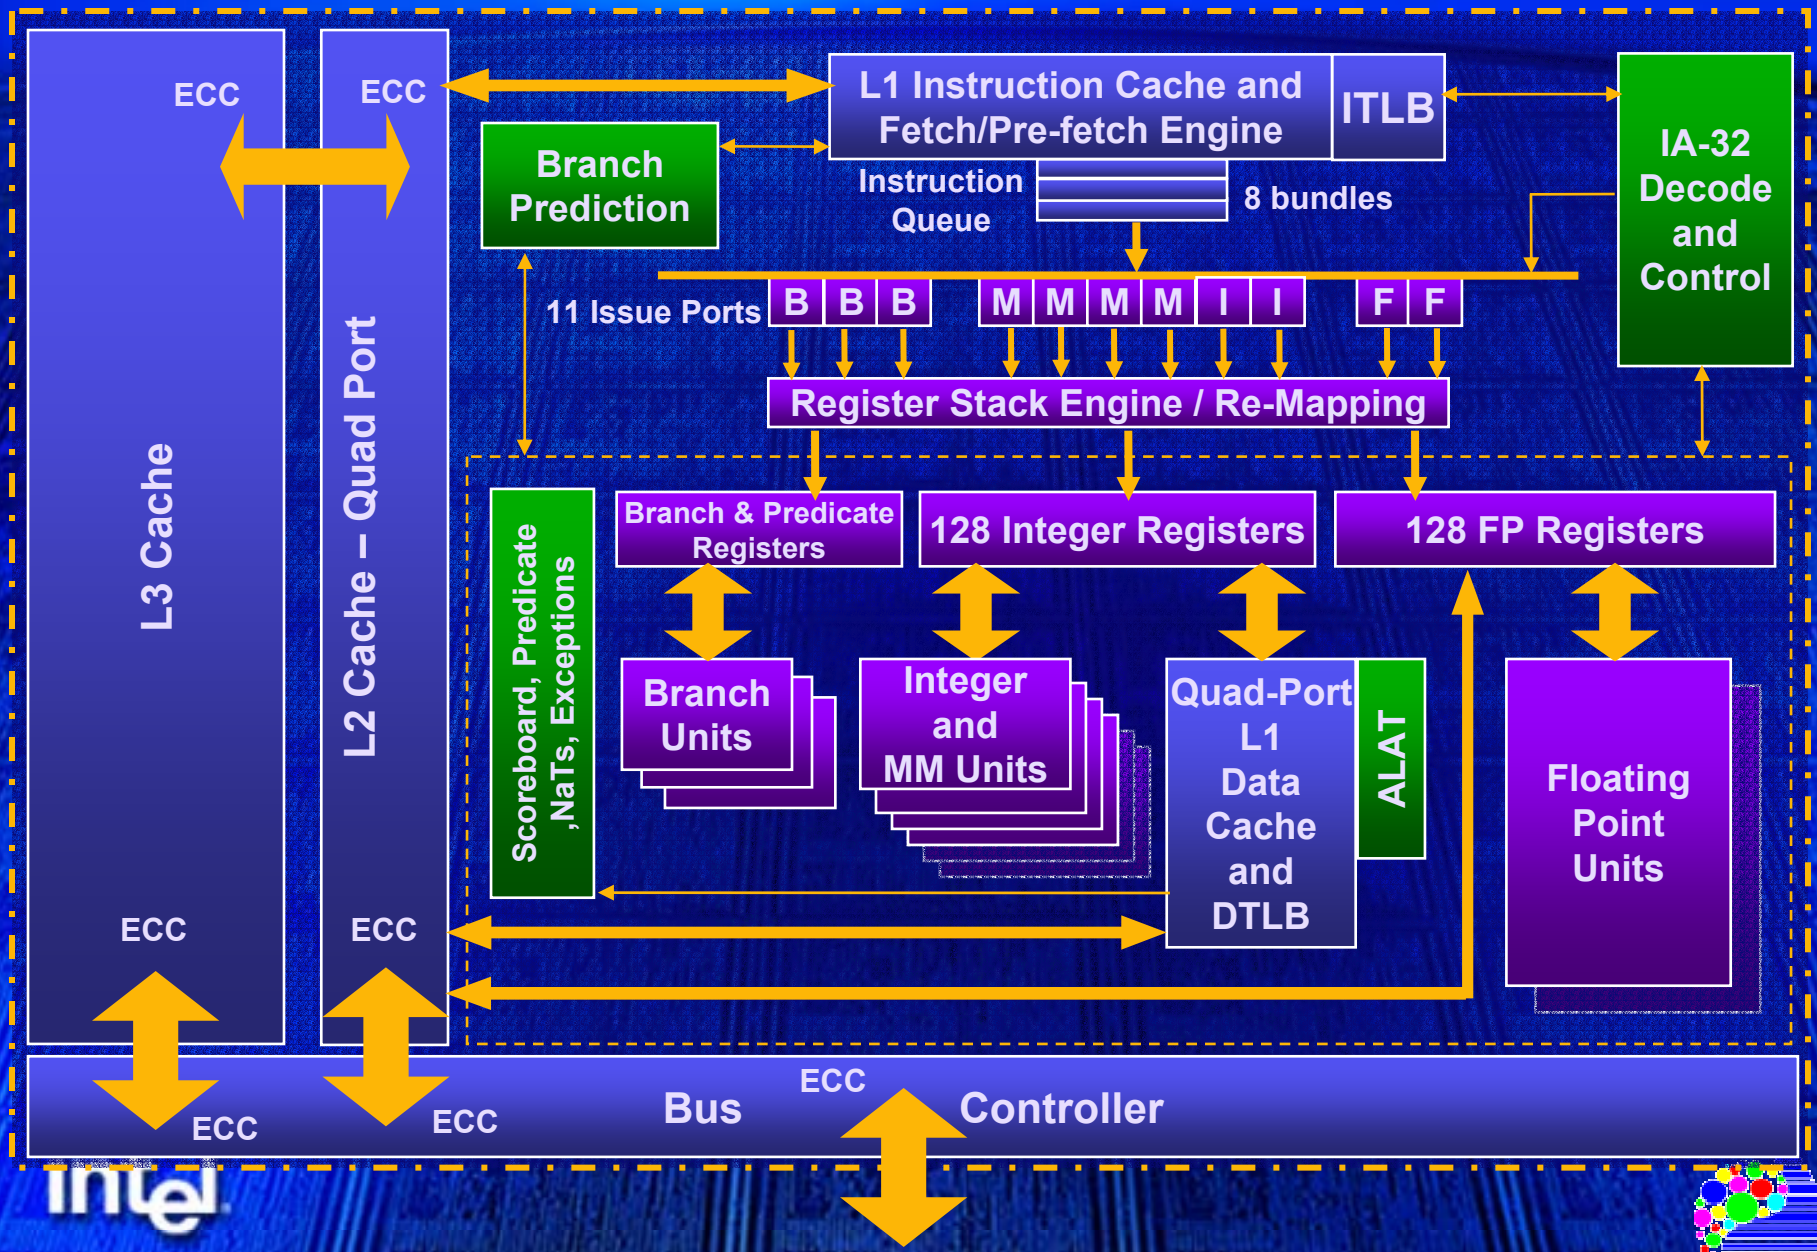

Integer & FP Performance

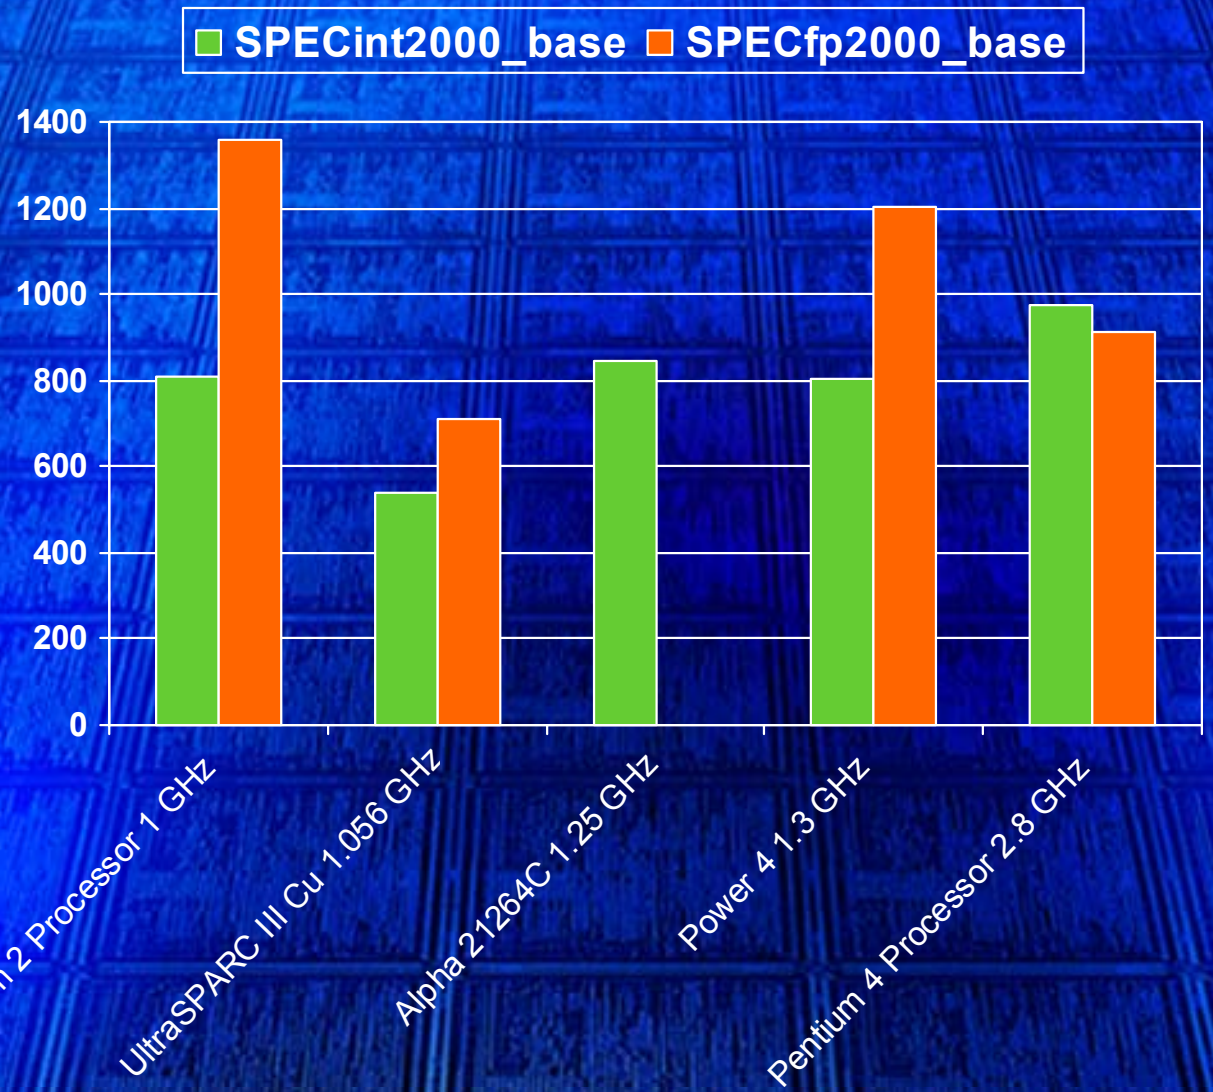

Source:www.specbench.org, 9/10/2002

Long Latency DRAM Accesses: Needs Memory Level Parallelism (MLP)

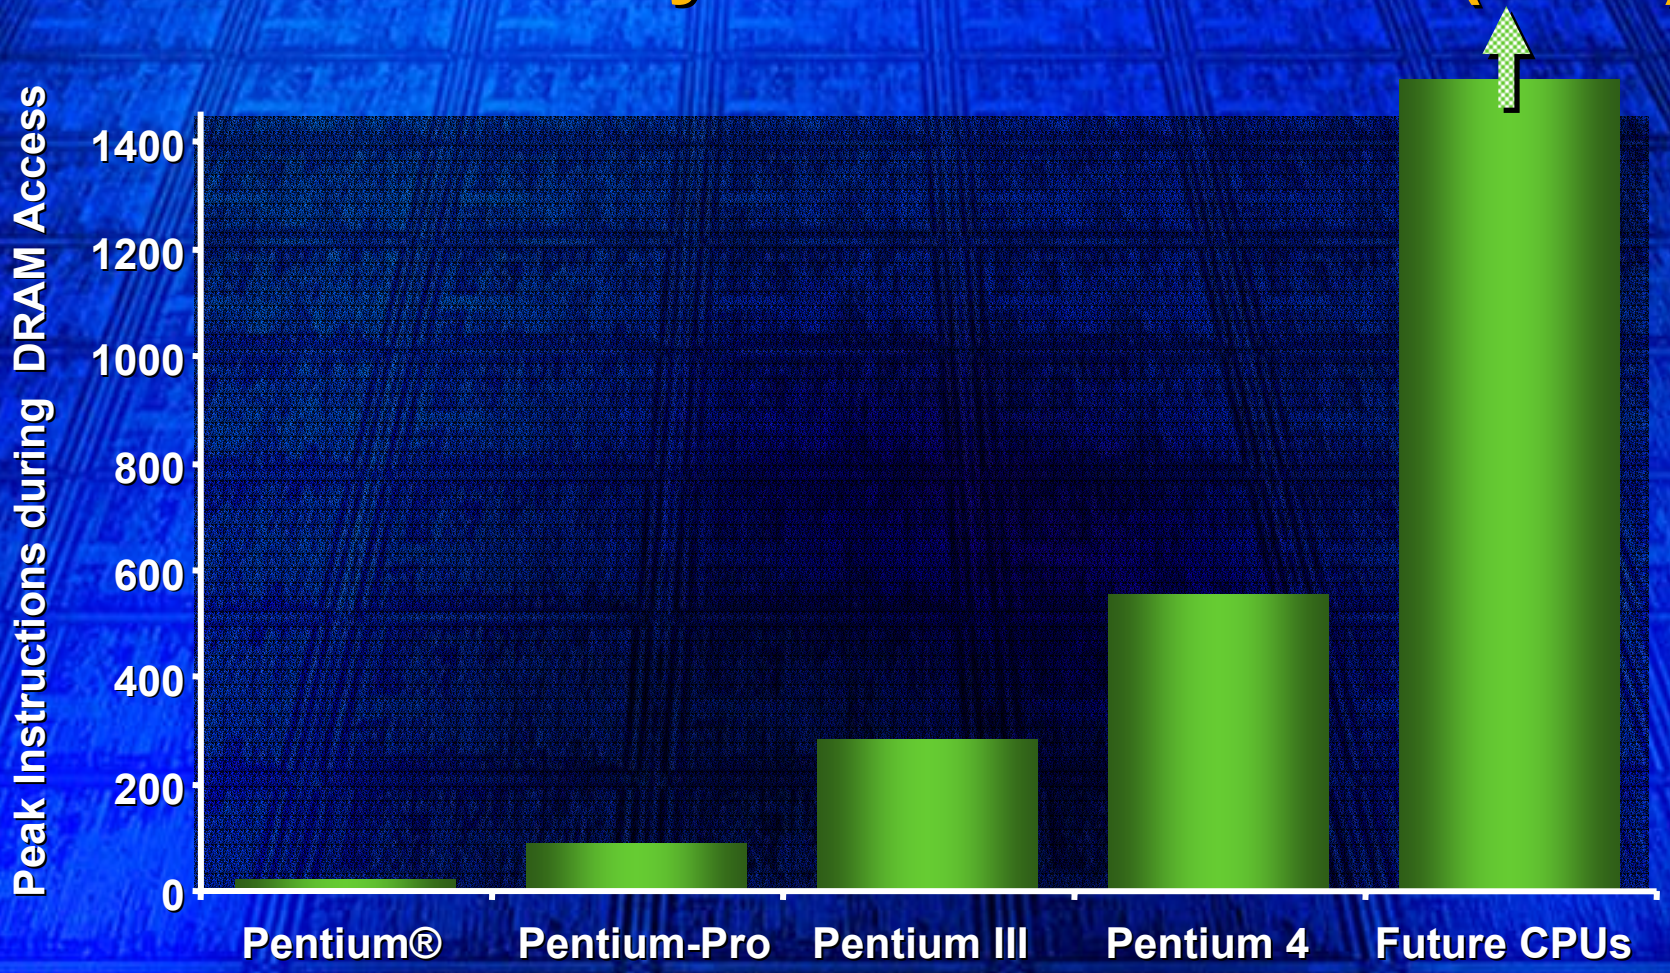

Multithreading

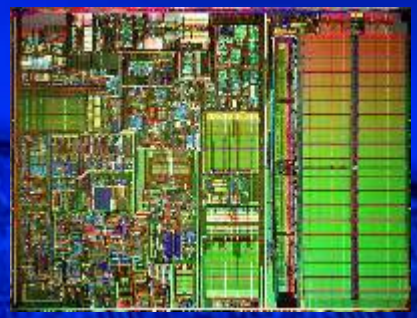

- Introduced on Intel® Xeon™ Processor MP
- Two logical processors for < 5% additional die area
- Executes two tasks simultaneously
  - Two different applications
  - Two threads of same application
- CPU maintains architecture state for two processors
  - Two logical processors per physical processor
- Power efficient performance gain
- 20-30% performance improvement on many throughput oriented workloads

IBM Power4 Dual Processor on a Chip

Two cores (~30M transistors each)

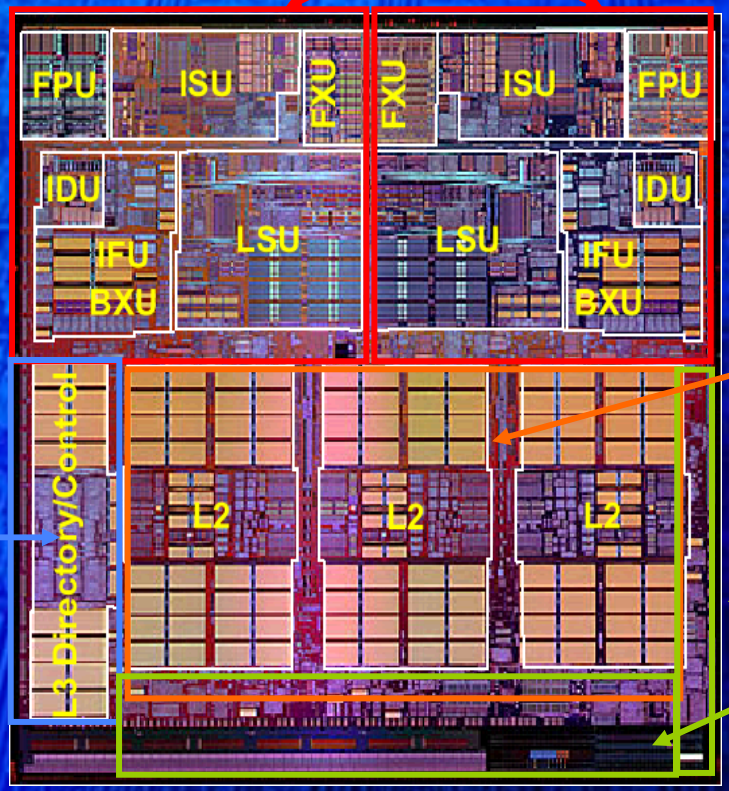

L3 & Mem Controller: L3 tags on-die for full-speed coherency checks

Large Shared L2: Multi-ported: 3 independent slices

Chip-to-Chip & MCM-to-MCM Fabric: Glueless SMP

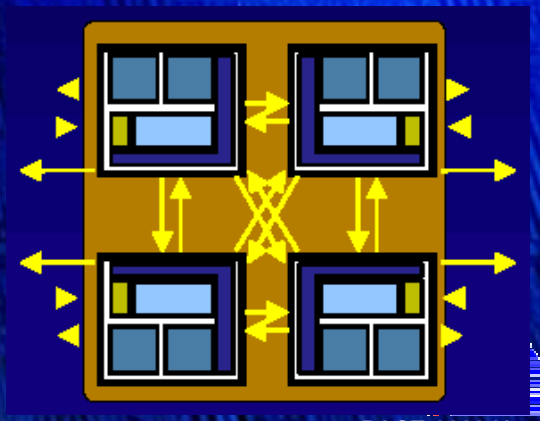

intel

*Other names and brands may be claimed as the property of others

HP PA-8800 Dual Processor on a Chip

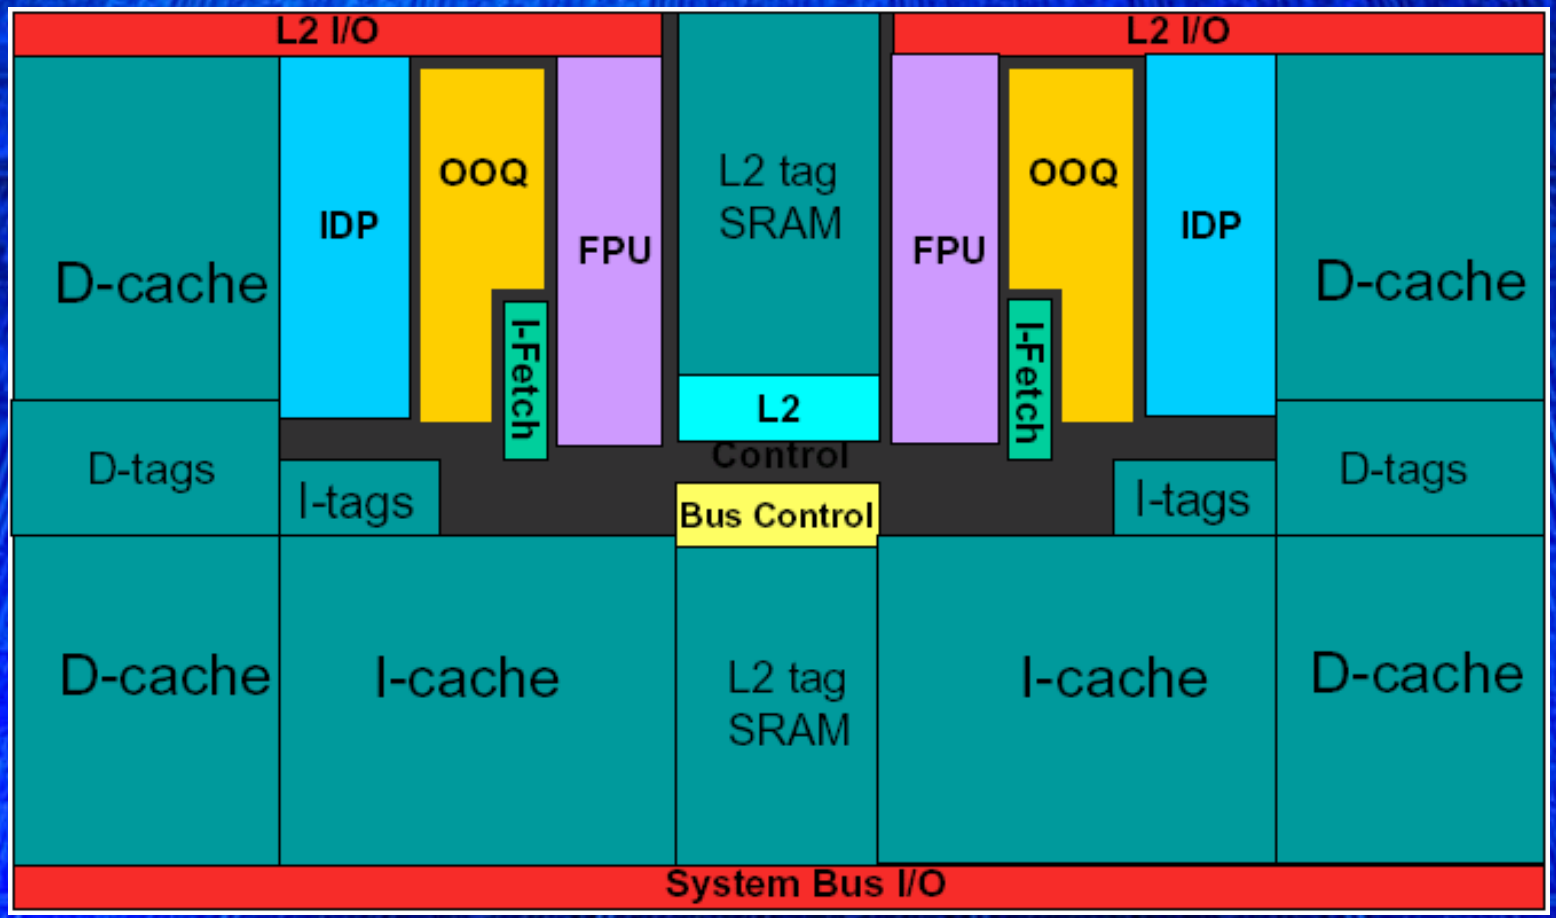

Outline

- Semiconductor Technology Evolution
- Moore's Law Video
- Parallelism in Microprocessors Today
- Multiprocessor Systems
- The Billion Transistor Chip
- Summary

Large Multiprocessor Systems

- 32 and 64 processor systems available today
- 300K to 400K transactions per minute
- >100 Linpack Gigaflops

IBM eServer pSeries 690

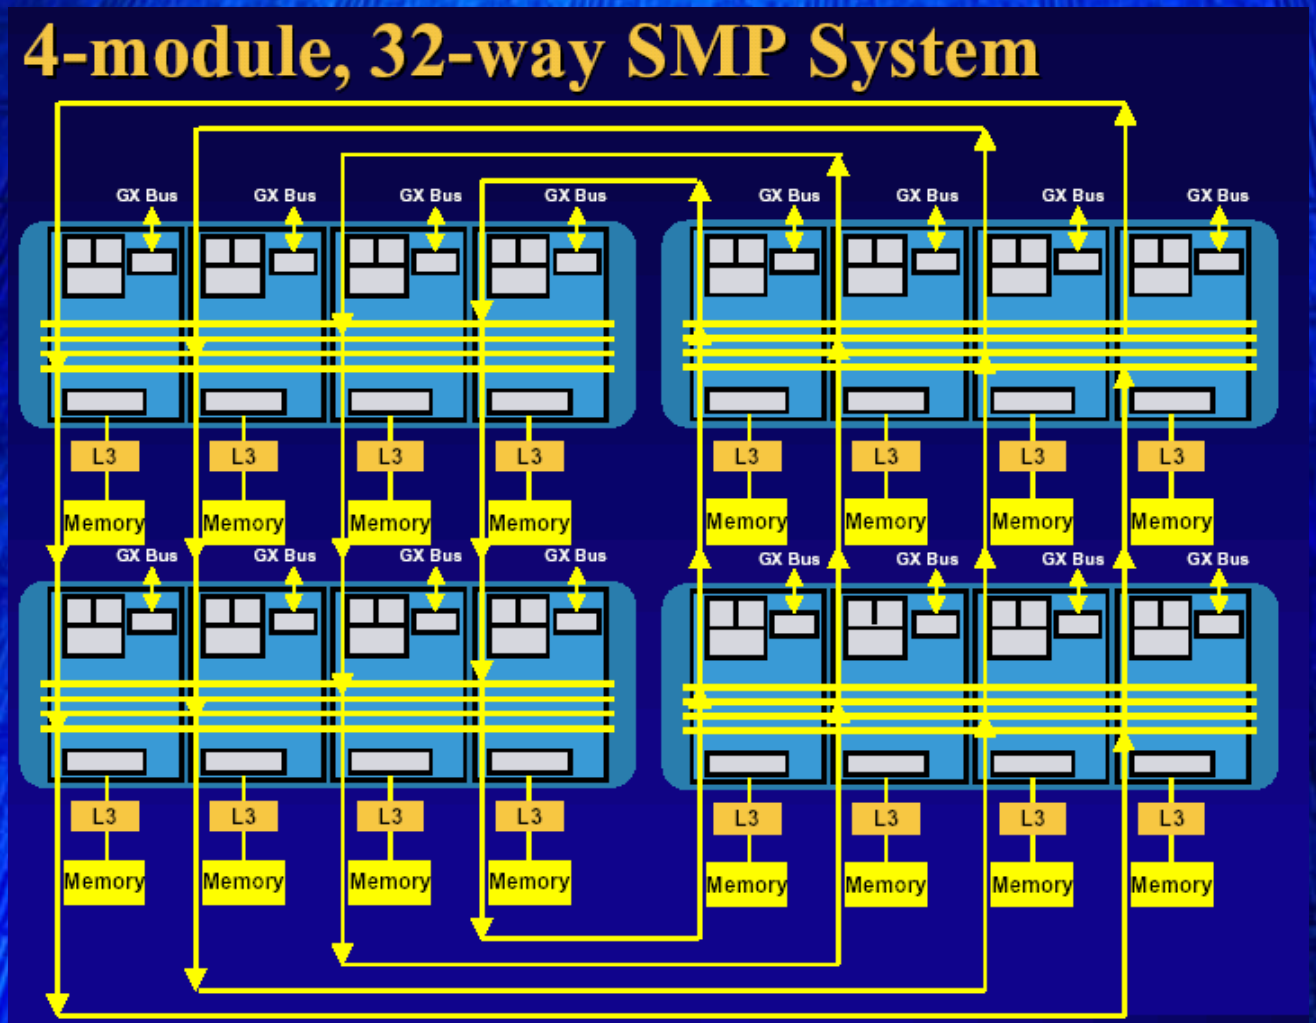

- 1.3 GHz Power4
- 8 to 32 CPU
- Starting at $450,000
- 8-way MCM @ $275,000**
- 403,255 tpmc @ $17.80 per tpmC
- 95 Linpack Gflops

***http://www-1024.ibm.com/content/home/store_IBMPublicUSA/en_US/eServer/pSeries/high_end/pSeries_highend.html
**Source: http://www.tpc.org/results/individual_results/IBM/IBMp690es_08142002.pdf
*Other names and brands may be claimed as the property of others

NEC TX7/i9510 SMP Server

- Up to 32 Itanium® 2 processors
- Up to 512GB memory (with 2GB DIMMs)
- Up to 112 PCI-X I/O slots
- Low latency and high bandwidth cross-bar interconnect
- Inter-cell memory interleaving
- ECC protected data transfer
- 308,620 tpmC @ $14.96 per tpmC
- 101 Linpack GigaFlops
- 32 Processors + 256GB @ $1,397,152**

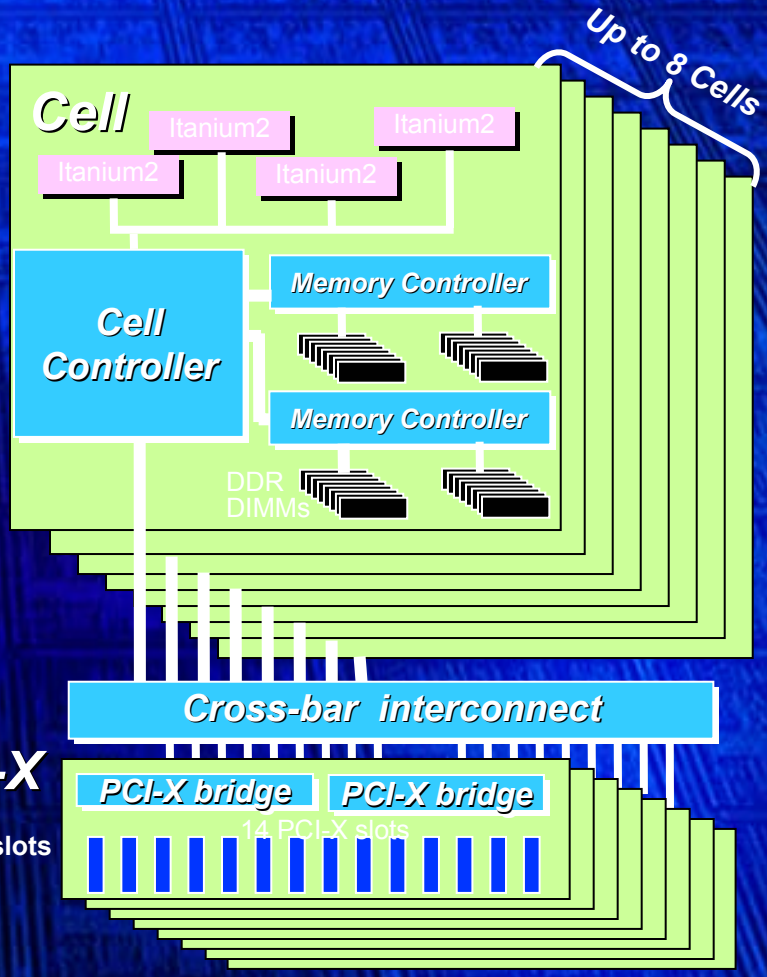

**http://www.tpc.org/results/individual_results/NEC/nec.tx7.i9510.c5.020909.es.pdf

*Other names and brands may be claimed as the property of others

HP Superdome

Super Dome is a cell-based hierarchical cross-bar system.

A cell consists of

- ➔ 4 CPUs
- ➔ 2 to 16GBs of Memory
- ➔ A link to 12 PCI I/O Slots
- ➔ Cell Board with 4 PA-8700 875MHz Processors @ $10.080** (2 chassis @ $424,275**)

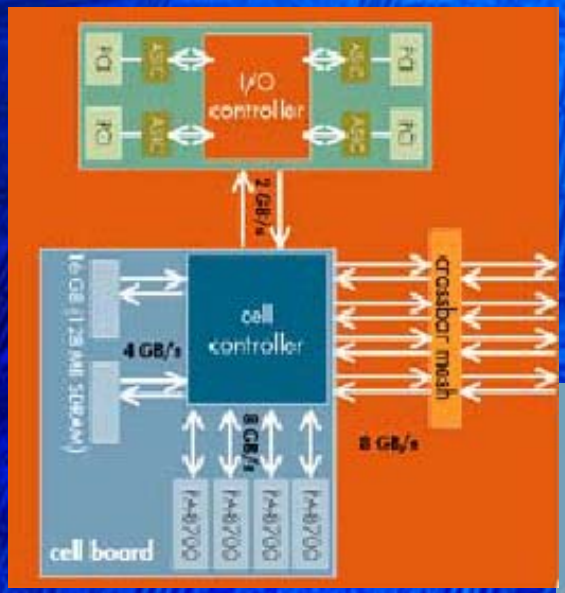

64P Performance

875 MHz PA-RISC 8700

- **423,414 tpmC @ $15.64 per tpmC**
- **134 Linpack Gigaflops**

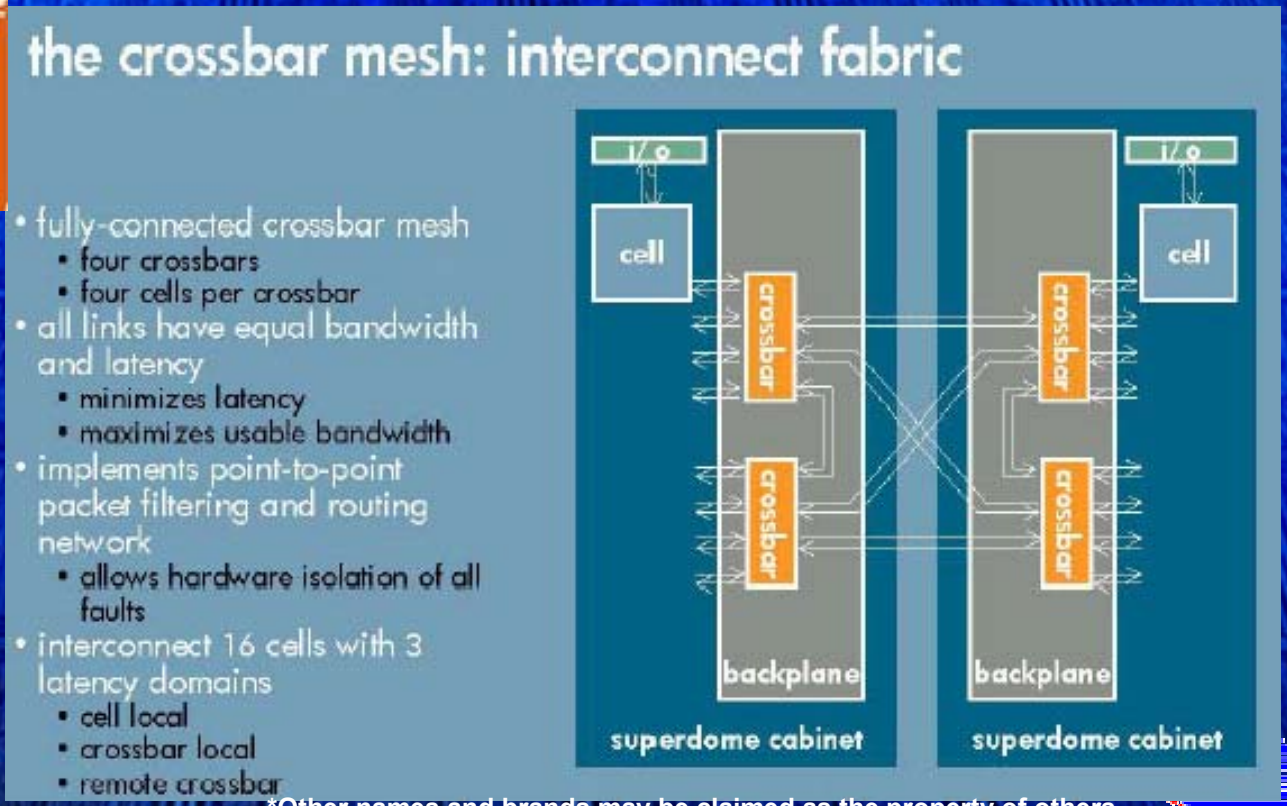

**http://www.tpc.org/results/individual_results/HP/hp_tpcc_sd_es.pd

*Other names and brands may be claimed as the property of others

HPC Clusters

15 node dual 2.4GHz Pentium® Xeon® cluster

- Commercial Off The Shelf (COTS) components
  - Processors
  - Packaging
  - Interconnects
  - Operating systems

3,000 Compaq Alpha EV68 microprocessors, housed in 750 four-processor AlphaServer systems running Tru64 UNIX

intel

Outline

- Semiconductor Technology Evolution
- Moore's Law Video
- Parallelism in Microprocessors Today
- Multiprocessor Systems
- The Billion Transistor Chip
- Summary

Parallelism Design Space

With Each Process Generation

- Frequency increases by about 1.5X
- Vcc will scale by only ~0.8
- Active power will scale by ~0.9
- Active power density will increase by ~30-80%
- Leakage power will make it even worse

Doubling performance requires more than 4 times the transistors

Single-stream Performance vs Costs

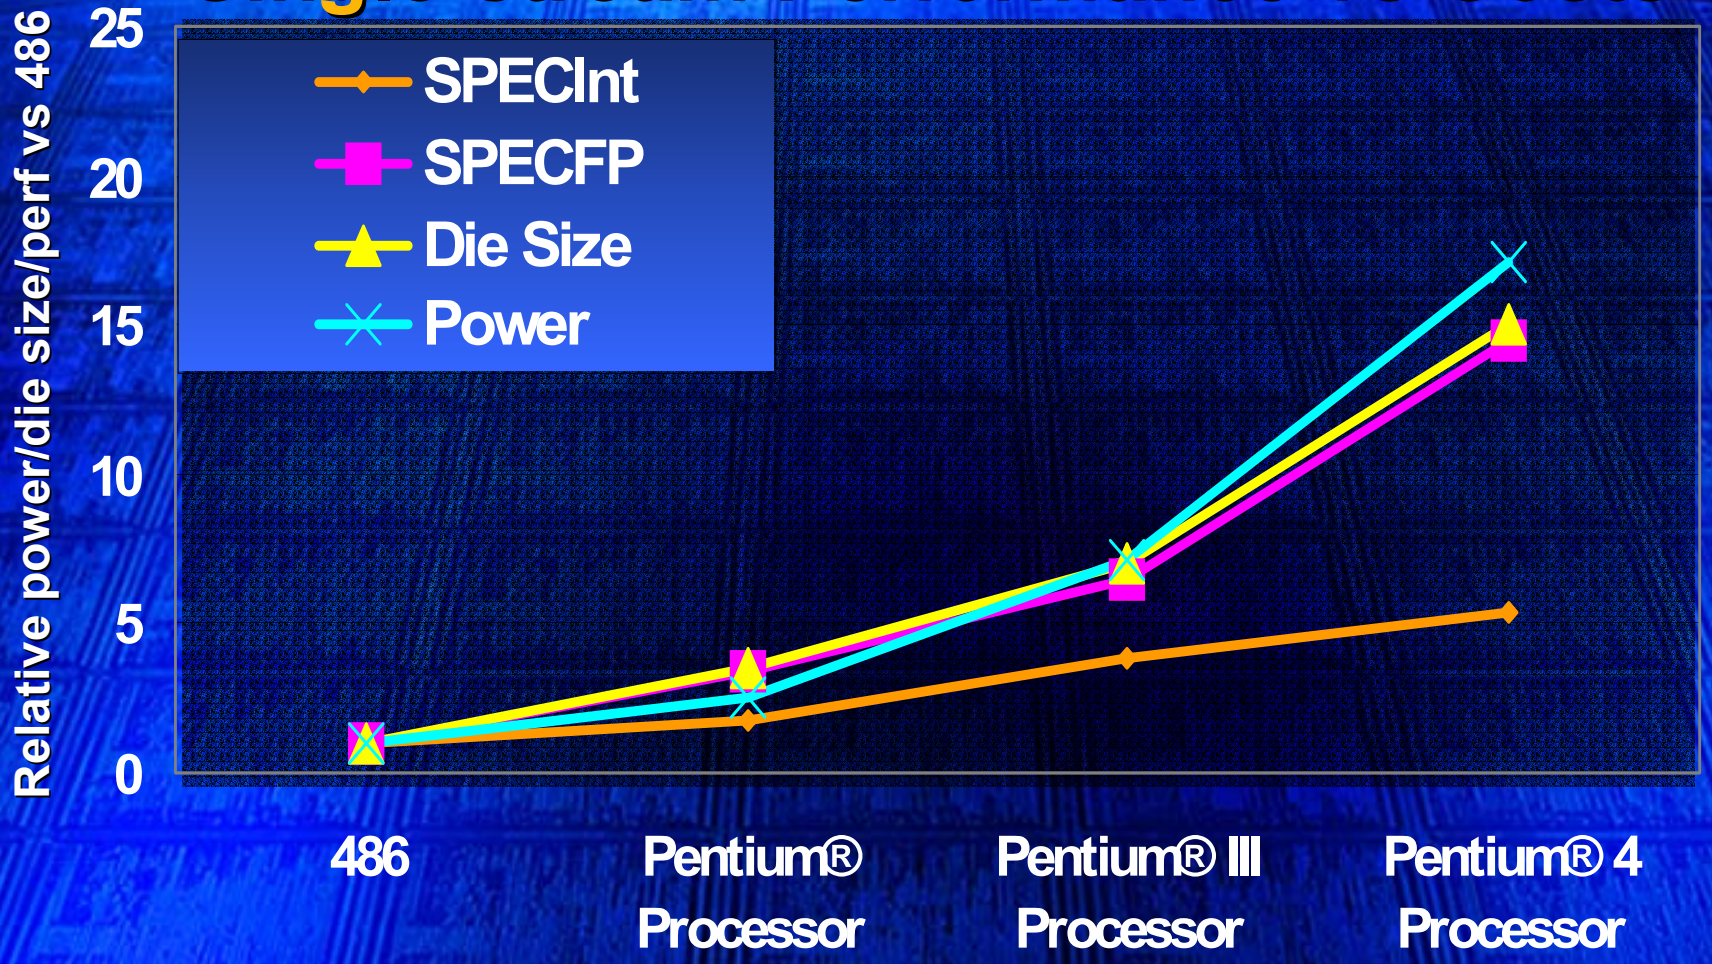

1999 Mainstream Microprocessor

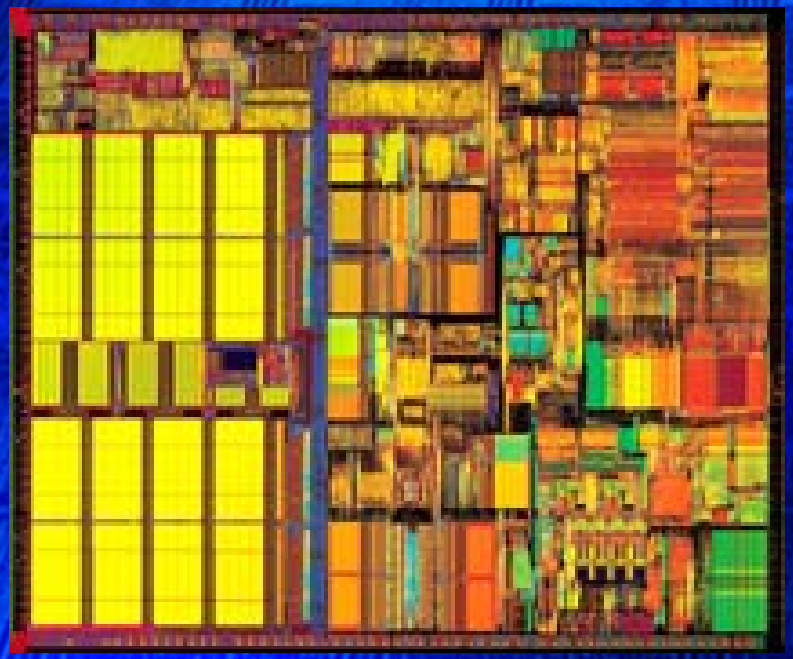

- Pentium® III Processor
- Integrated 256 KB L2 cache
- 106 mm² die size
- 0.18µ process
- 6 metal layer process
- 28 million transistors

Technology Projection

	1999	2001	2003	2005	2007
Process	180 nm	130 nm	90 nm	65 nm	50 nm
Core Sq mm	50-100	25-50	12-25	6-12	3-6
1 MB cache Sq mm	100	50	25	12	6
# of cores in ~ 200 sq mm	2-4	4-8	8-16	16-32	32-64
MB of cache in ~200 sq mm	~2	~4	~8	~16	~24

Art of the Possible

- Billion Transistors possible in 2005
- Large die sizes can be built
  - 4 to 6 square centimeters
- What can fit on a single die in 2005?
  - 12 mm^2 per processor
  - 12 mm^2 per MB

Die size in mm^2	4 cores	8 cores	16 cores
16 MB cache	240	288	384
32 MB cache	432	480	576

CMP Challenges

- How much Thread Level Parallelism is there in non-embarassingly parallel workloads?
- Ability to generate code with lots of threads
- Thread synchronization
- Operating systems for parallel machines
- Single thread performance
- Power limitations
- On-chip interconnect infrastructure

Outline

- Semiconductor Technology Evolution
- Moore's Law Video
- Parallelism in Microprocessors Today
- Multiprocessor Systems
- The Billion Transistor Chip
- Summary

Summary

- Plenty of opportunities for “parallel programming” in Commercial Off The Shelf Server platforms
- Amount of parallelism in hardware will increase
- Need applications and tools that can exploit parallelism at all levels