

Microprocessor Microarchitecture

Introduction



Lynn Choi

School of Electrical Engineering



高麗大學校

Computer System Laboratory

Class Information



❑ *Lecturer*

- Prof. Lynn Choi, 02-3290-3249, lchoi@korea.ac.kr

❑ *Textbook*

- Computer Architecture, A Quantitative Approach
 - 5th edition, Hennessy and Patterson, Morgan Kaufmann
- Lecture slides (collection of research papers)
- Reading list (refer to the class homepage)

❑ *Content*

- Introduction
- Branch Prediction
- Instruction Fetch
- Data Hazard and Dynamic Scheduling
- Limits on ILP
- Exceptions
- Multiprocessors and Multithreading
- Advanced Cache Design and Memory Hierarchy
- IA64 and Itanium CPU

Class Information



❑ *Special Topics*

- Multicore and manycore processors
- Presentation of ~2 papers in the subject

❑ *Project*

- Research proposal
- Simulation and experimentation results
- Detailed survey

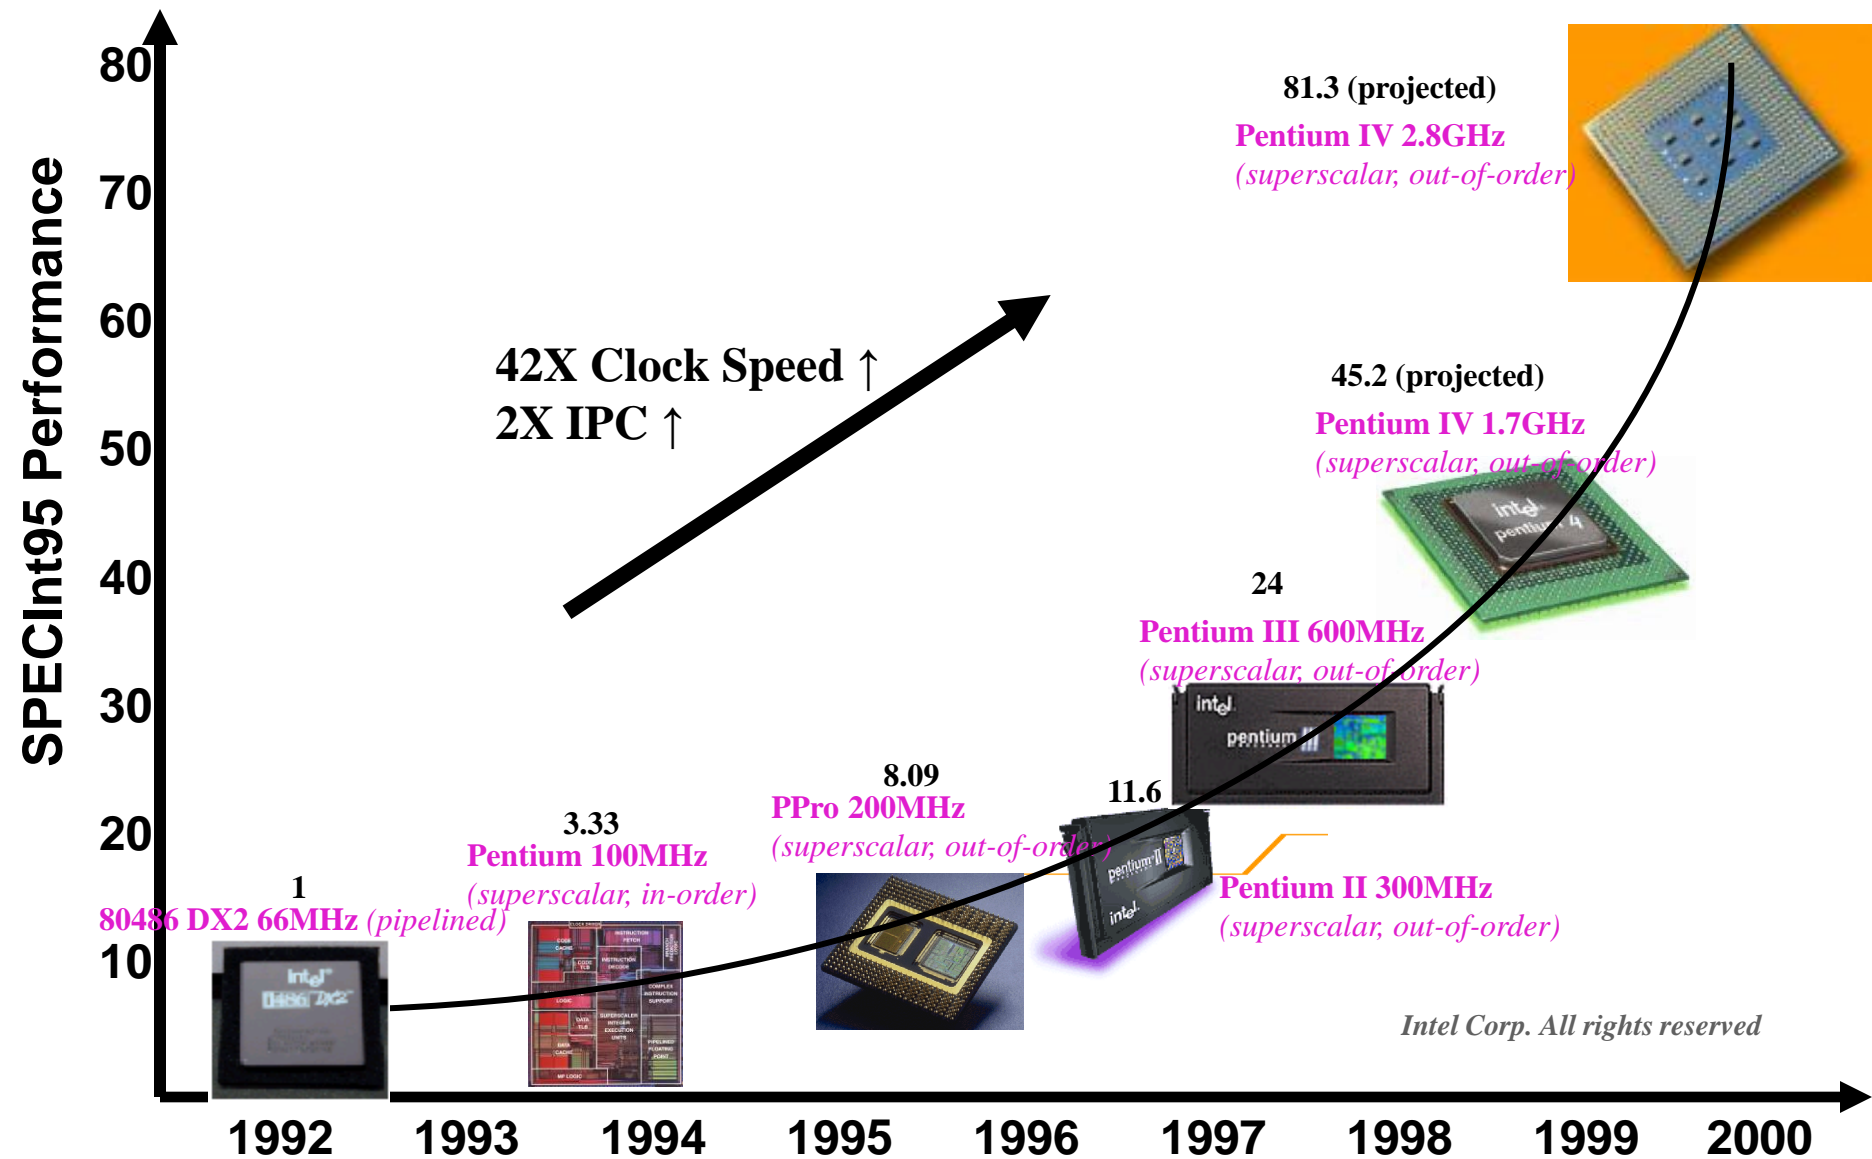
❑ *Evaluation*

- Midterm : 35%
- Final: 35%
- Presentation: 15%
- Project: 15%

❑ *Class organization*

- Lecture: 70%
- Presentation: 30% (after Midterm)

Advances in Intel Microprocessors



Intel Corp. All rights reserved

Intel® Pentium 4 Microprocessor



□ *Intel Pentium IV Processor*

Intel Corp. All rights reserved



- Technology
 - 0.13 μ process, 55M transistors, 82W
 - 3.2 GHz, 478pin Flip-Chip PGA2
- Performance
 - 1221 Ispec, 1252 Fspec on SPEC 2000
 - ▼ Relative performance to SUN 300MHz Ultra 5_10 workstation (100 Ispec/Fspec)
 - 40% higher clock rate, 10~20% lower IPC compared to P III
- Pipeline
 - 20-stage out-of-order (OOO) pipeline, hyperthreading
 - 2 ALUs run at 6.4GHz
- Cache hierarchy
 - 12K micro-op trace cache/8 KB on-chip D cache
 - On-chip 512KB L2 ATC (Advanced Transfer Cache)
 - Optional on-die 2MB L3 Cache
- 800MHz system bus, 6.4GB/s bandwidth
 - Implemented by quad-pumping on 200MHz system bus

Intel® Itanium® 2 processor




□ Intel® Itanium® 2 processor

Intel Corp. All rights reserved



- Technology
 - 1.5 GHz, 130W
- Performance: 1322 Ispec, 2119 Fspec
 - 50% higher transaction performance compared to Sun UltraSPARC III Cu processor (4-way MP system)
- EPIC architecture
- Pipeline
 - 8-stage in-order pipeline (10-stage in Itanium)
 - 11 issue ports (9 ports in Itanium)
 - 6 INT, 4 MEM, 2 FP, 1 SIMD, 3 BR (4 INT, 2 MEM in Itanium)
- Cache hierarchy
 - 32KB L1 cache, 256KB L2 cache, and up to 6MB L3 Cache
- Memory and System Interface
 - 50b PA, 64b VA
 - 400MHz 128-bit system bus, 6.4GB/s bandwidth (compared to 266MHz 64-bit system bus, 2.1GB.s in Itanium)

Microprocessor Performance Curve



Today's Microprocessor

❑ Intel i7 Processor

➤ Technology

- 32nm process, 130W, 239 mm² die
- 3.46 GHz, 64-bit 6-core 12-thread processor
- 159 Ispec, 103 Fspec on SPEC CPU 2006 (296MHz UltraSparc II processor as a reference machine)

➤ Core microarchitecture

- Next generation *multi-core* microarchitecture introduced in Q1 2006 (Derived from P6 microarchitecture)
- Optimized for multi-cores and lower power consumption
 - 14-stage 4-issue out-of-order (OOO) pipeline
- 64bit Intel architecture (x86-64)
- Core i3 (entry-level), Core i5 (mainstream consumer), Core i7 (high-end consumer), Xeon (server)

➤ 256KB L2 cache/core, 12MB L3 Caches

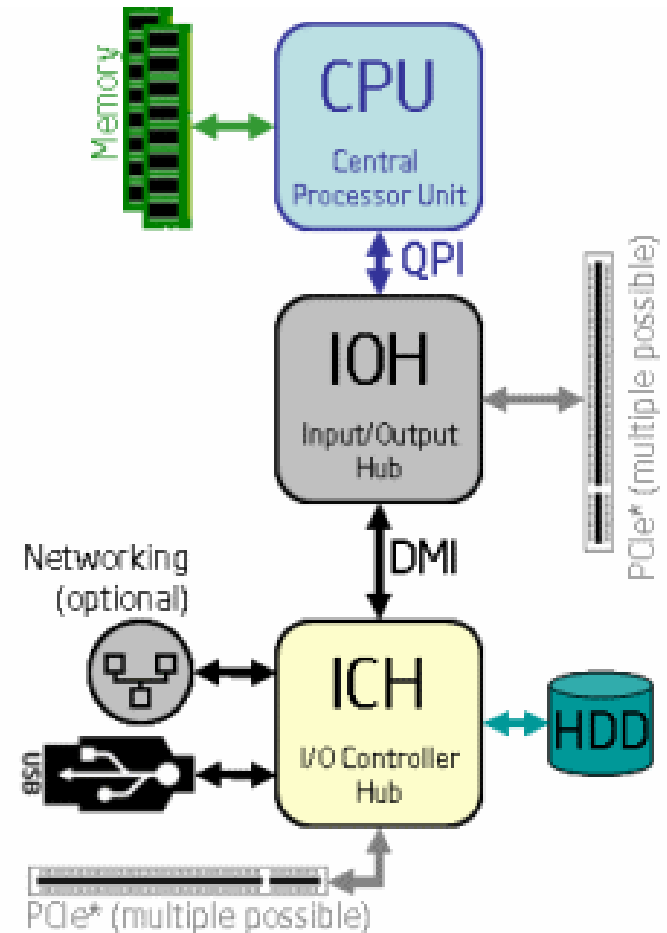
➤ Integrated memory controller



Intel Corp. All rights reserved

Intel i7 System Architecture

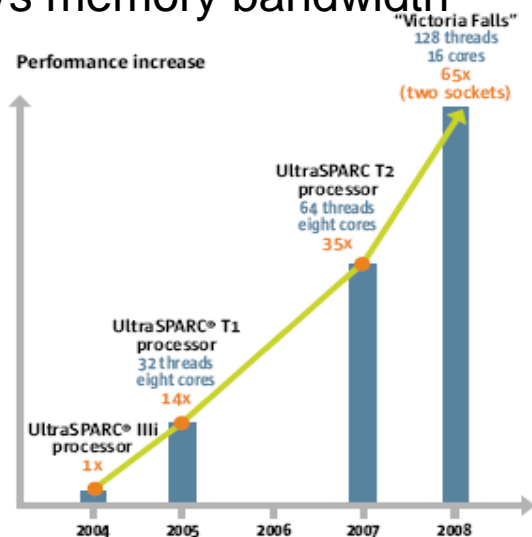
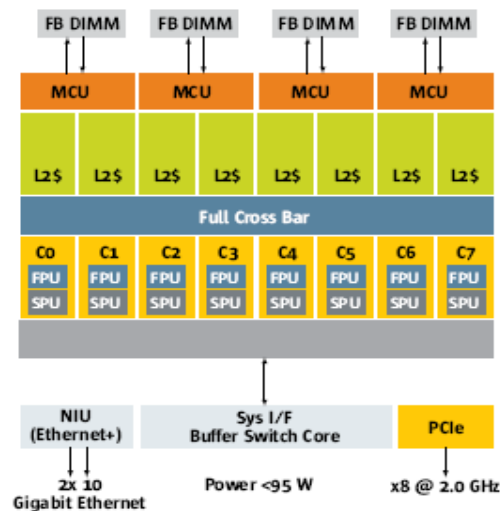
- ❑ **Integrated memory controller**
 - 3 Channel, 3.2GHz clock, 25.6 GB/s memory bandwidth (memory up to 24GB DDR3 SDRAM), 36 bit physical address
- ❑ **QuickPath Interconnect (QPI)**
 - Point-to-point processor interconnect, replacing the front side bus (FSB)
 - 64bit data every two clock cycles, up to 25.6GB/s, which doubles the theoretical bandwidth of 1600MHz FSB
- ❑ **Direct Media Interface (DMI)**
 - The link between Intel Northbridge and Intel Southbridge, sharing many characteristics with PCI-Express
- ❑ **IOH (Northbridge)**
- ❑ **ICH (Southbridge)**



Today's Microprocessor

□ Sun UltraSPARC T2 processor (“Niagara II”)

- Multithreaded multicore technology
 - Eight 1.4 GHz cores, 8 threads per core → total 64 threads
 - 65nm process, 1831 pin BGA, 503M transistors, 84W power consumption
- Core microarchitecture: Two issue 8-stage instruction pipelines
- 4MB L2 – 8 banks, 64 FB DIMMs, 60+ GB/s memory bandwidth



Oracle. All rights reserved

□ Sun UltraSPARC T3 processor (“Rainbow Falls”)

- 40nm process, 16 1.65GHz cores, 8 threads per core → total 128 threads

Trends in Technology

- **Integrated circuit technology**
 - Transistor density: 35%/year
 - Die size: 10-20%/year
 - Integration overall: 40-55%/year

- **DRAM capacity: 25-40%/year (slowing)**

- **Flash capacity: 50-60%/year**
 - 15-20X cheaper/bit than DRAM

- **Magnetic disk technology: 40%/year**
 - 15-25X cheaper/bit than Flash
 - 300-500X cheaper/bit than DRAM

Bandwidth and Latency

□ Bandwidth or throughput

- Total work done in a given time
- 10,000-25,000X improvement for processors
- 300-1200X improvement for memory and disks

□ Latency or response time


- Time between start and completion of an event
- 30-80X improvement for processors
- 6-8X improvement for memory and disks

□ Feature size

- Minimum size of transistor or wire in x or y dimension
- 10 microns in 1971 to .032 microns in 2011
- Transistor performance scales linearly
- Integration density scales (*more than*) quadratically
- However, wire delay scales poorly compared to transistor performance!
 - In the past few years, both wire delay and power dissipation have become major design limitations for VLSI design

Copyright © 2012, Elsevier Inc. All rights reserved.

Bandwidth and Latency



Log-log plot of bandwidth and latency milestones

Copyright © 2012, Elsevier Inc. All rights reserved.

Dynamic Power



- ❑ *For CMOS chips, traditional dominant energy consumption has been in switching transistors, called **dynamic power***

$$Power_{dynamic} = 1/2 \times CapacitiveLoad \times Voltage^2 \times FrequencySwitched$$

- For a fixed task, slowing clock rate (frequency switched) reduces power, but not energy
 - Dropping voltage helps both, so went from 5V to 1V
 - Capacitive load is a function of number of transistors connected to output and technology determines capacitance of wires and transistors
- ❑ *To save energy & dynamic power, most CPUs now turn off clock of inactive modules (e.g. FPU)*

Example



- *Suppose 15% reduction in voltage results in a 15% reduction in frequency. What is impact on dynamic power?*

$$\begin{aligned} Power_{dynamic} &= 1/2 \times CapacitiveLoad \times Voltage^2 \times FrequencySwitched \\ &= 1/2 \times CapacitiveLoad \times (.85 \times Voltage)^2 \times .85 \times FrequencySwitched \\ &= (.85)^3 \times OldPower_{dynamic} \\ &\approx 0.6 \times OldPower_{dynamic} \end{aligned}$$

Static Power



- ❑ *Because leakage current flows even when a transistor is off, static power important too*

$$Power_{static} = Current_{static} \times Voltage$$

- ❑ *Leakage current increases in processors with smaller transistor sizes*
- ❑ *In 2006, goal for leakage is 25% of total power consumption; high performance designs at 40%*
- ❑ *Very low power systems even gate voltage to inactive modules to control loss due to leakage*

Processor Performance Equation



□ T_{exe} (*Execution time per program*)


$$= NI * CPI_{execution} * T_{cycle}$$

- NI: # of instructions / program (program size)
 - Small program is better
- CPI: clock cycles / instruction
 - Small CPI is better. In other words, higher IPC is better
- T_{cycle} = clock cycle time
 - Small clock cycle time is better. In other words, higher clock speed is better

Clock Speed versus Power



- ❑ Intel 80386 consumed ~ 2 W
- ❑ 3.3 GHz Intel Core i7 consumes 130 W
- ❑ Heat must be dissipated from 1.5 x 1.5 cm chip
- ❑ This is the limit of what can be cooled by air



Definition: Performance



$$\text{Performance}(x) = \frac{1}{\text{Execution_time}(x)}$$

"X is n times faster than Y" means

$$n = \frac{\text{Performance}(X)}{\text{Performance}(Y)} = \frac{\text{Execution_time}(Y)}{\text{Execution_time}(X)}$$

Performance: What to measure



- ❑ *Usually rely on benchmarks vs. real workloads*
- ❑ *To increase predictability, collections of benchmark applications, called **benchmark suites**, are popular*
- ❑ ***SPECCPU: popular desktop benchmark suite***
 - CPU only, split between integer and floating point programs
 - SPECint2000 has 12 integer, SPECfp2000 has 14 FP programs
 - SPECCPU2006 is announced Spring 2006
 - 12 integer and 17 FP programs
- ❑ ***Transaction Processing Council measures server performance and cost-performance for databases***
 - TPC-C Complex query for Online Transaction Processing
 - TPC-H models ad hoc decision support
 - TPC-W a transactional web benchmark
 - TPC-App application server and web services benchmark

SPEC Benchmark Evolution



SPEC2006 benchmark description	Benchmark name by SPEC generation				
	SPEC2006	SPEC2000	SPEC95	SPEC92	SPEC89
GNU C compiler					gcc
Interpreted string processing			perl		espresso
Combinatorial optimization		mcf			li
Block-sorting compression		bzip2		compress	eqntott
Go game (AI)	go	vortex	go	sc	
Video compression	h264avc	gzip	ijpeg		
Games/path finding	astar	eon	m88ksim		
Search gene sequence	hmmer	twolf			
Quantum computer simulation	libquantum	vortex			
Discrete event simulation library	omnetpp	vpr			
Chess game (AI)	sjeng	crafty			
XML parsing	xalancbmk	parser			
CFD/blast waves	bwaves				fpppp
Numerical relativity	cactusADM				tomcatv
Finite element code	calculix				doduc
Differential equation solver framework	dealll				nasa7
Quantum chemistry	gamess				spice
EM solver (freq/time domain)	GemsFDTD			swim	matrix300
Scalable molecular dynamics (~NAMD)	gromacs		apsi	hydro2d	
Lattice Boltzman method (fluid/air flow)	lbm		mgrid	su2cor	
Large eddie simulation/turbulent CFD	LESlie3d	wupwise	applu	wave5	
Lattice quantum chromodynamics	milc	apply	turb3d		
Molecular dynamics	namd	galgel			
Image ray tracing	povray	mesa			
Spare linear algebra	soplex	art			
Speech recognition	sphinx3	equake			
Quantum chemistry/object oriented	tonto	facerec			
Weather research and forecasting	wrf	ammp			
Magneto hydrodynamics (astrophysics)	zeusmp	lucas			
		fma3d			
		sixtrack			

© 2007 Elsevier, Inc. All rights reserved.

How Summarize Suite Performance (1/3)



- ❑ *Arithmetic average of execution time of all programs?*
 - But they vary by 4X in speed, so some would be more important than others in arithmetic average
- ❑ *Could add a weights per program, but how pick weight?*
 - Different companies want different weights for their products
- ❑ **SPECRatio**: *Normalize execution times to reference computer, yielding a ratio proportional to performance*
=
$$\frac{\text{time on reference computer}}{\text{time on computer being rated}}$$

How Summarize Suite Performance (2/3)



- *If SPECRatio on Computer A is 1.25 times bigger than Computer B, then*

$$\begin{aligned} 1.25 &= \frac{SPECRatio_A}{SPECRatio_B} = \frac{\frac{ExecutionTime_{reference}}{ExecutionTime_A}}{\frac{ExecutionTime_{reference}}{ExecutionTime_B}} \\ &= \frac{ExecutionTime_B}{ExecutionTime_A} = \frac{Performance_A}{Performance_B} \end{aligned}$$

- *Note that when comparing 2 computers as a ratio, execution times on the reference computer drop out, so choice of reference computer is irrelevant*

How Summarize Suite Performance (3/3)



- *Since we use ratios, proper mean is geometric mean (SPECRatio unitless, so arithmetic mean meaningless)*

$$\textit{GeometricMean} = \sqrt[n]{\prod_{i=1}^n \textit{SPECRatio}_i}$$

Exercises & Discussion



- ❑ *3.2GHz Pentium4 processor is reported to have SPECint ratio of 1221 and SPECfp ratio of 1252 in SPEC2000 benchmarks. What does this mean?*
- ❑ *How much memory can you address using 38 bits of address assuming byte-addressability?*
- ❑ *Classify Intel's 32bit microprocessors in terms of processor generations from 80386 to Pentium 4. What's the meaning of generation here?*
- ❑ *Assume two processors, one RISC and one CISC implemented at the same clock speed and the same IPC. Which one performs better?*

Homework 1



- ❑ ***Read Chapter 1 and Chapter 2***
- ❑ ***Exercise***
 - 1.4
 - 1.5
 - 1.10
 - 1.13
 - 1.18