Architecture des Ordinateurs II

—

Part III: Case Studies
IA-32 and Pentium

Paolo.Ienne@epfl.ch

EPFL – I&C – LAP

Eduardo.Sanchez@epfl.ch

EPFL – I&C – LSL

ÉCOLE POLYTECHNIQUE FÉDÉRALE DE LAUSANNE

Intel Processors

Processor	Date	f (MHz)	Trans.	Features
4004	4/71	0.108	2300	First µP
8008	4/72	0.108	3500	First 8-bit µP
8080	4/74	2	6000	Popular 8-bit
8086	6/78	5-10	29k	First 16-bit µP; 20-bit addressing
8088	6/79	5-8	29k	Simpler; IBM PC
80286	2/82	8-12	134k	Protected mode, 24-bit addressing
80386	10/85	16-33	275k	32-bit (**IA-32**)
80486	4/89	25-100	1.2M	Pipelined (5-stage); cache
Pentium	3/93	60-233	3.1M	Superscalar, dual pipeline
PentiumPro	3/95	150-200	5.5M	Out-of-order; L2 cache
Pentium II	5/97	233-400	7.5M	MMX (SIMD instructions)
Pentium III	3/99	450-1200	9.5-26M	SSE (incl. SIMD-FP); 10-stage pipeline
Pentium 4	12/00	1300-2200	42M	SSE2 (128-bit); TC; 20-stage pipeline

Intel Processors

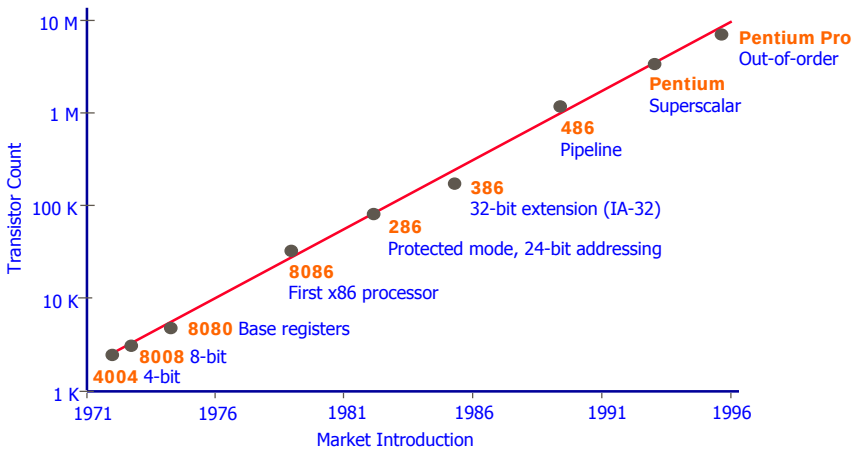

Current High-End Processors

Source: Microprocessor Report, © Cahners 2006

Processor	Alpha 21364 EV-78+	AMD Opteron 254	AMD Dual-core Opteron 280	HP PA-8900	IBM Power4+	IBM Power5
Processor Arch	64-bit	32/64-bit	Dual 32/64-bit	Dual 64-bit	Dual 64-bit	Dual, MT 64-bit
Clock Rate	1.30GHz	2.8GHz	2.4GHz	1.16GHz	1.7GHz	1.9GHz
Cache (I/D/L2/L3)	64K/64K/ 1.75M	64K/64K/ 1M	2 x 64K/64K/1M	1.5M/1.5M/ 64M	64K/32K/ 1.5MB	64K/32K/ 1.92MB/36MB
Issue Rate/Core	4 issue	3 x86 instr	3 x86 instr	4 issue	8 issue	8 issue
Pipeline Stages	7/9 stages	9/11 stages	9/11 stages	7/9 stages	12/17 stages	12/17 stages
Out of Order	80 instr	72ROPs	72ROPs	56 instr	200 instr	200 instr
Rename Regs	48/41	36/36	36/36	56 total	48/40	48/40
BHT Entries	4K x 9-bit	4K x 2-bit	4K x 2-bit	8K x 2-bit	3x16K x 1-bit	3x16K x 1-bit
TLB Entries	128/128	280/288	280/288	2 x 240 unified	2x1,024 unified	2x1,024 unified
Memory B/W	12GB/s	6.4GB/s	6.4GB/s	6.4GB/s	12.8GB/s	12.8GB/s
Package	FC-LGA1443	PGA-940	PGA-940	LGA-544	MCM	MCM
IC Process	0.18µm 7M	0.13µm 6M	0.09µm 7M	0.13µm 7M	0.13µm 7m	0.13µm 7m
Die Size	397mm²	193mm²	199mm²	304mm²	267mm²**	389mm²**
Transistors	135 million	106 million	233 million	300 million	184 million**	276 million**
Est Die Cost	$180	$79	$85	$96	$144**	$200**
Power (Max)	155W	92W(MTP)*	95W(MTP)	103W	100W**	120W*
Availability	3Q04	4Q05	4Q05	3Q03	2Q03	4Q05
Configuration	2–64 way	1–2 way	1–2 way	1–128 way	2–32 way	2–32 way
SPEC_int2000(base)	904	1,817	1,499	N/A	1,077	1,470
SPEC_fp2000(base)	1,279	2,132	1,752	N/A	1,598	2,839

Processor	Intel Itanium 2	Intel XeonMP	Intel Xeon	MIPS R16000	Fujitsu SPARC64 V	Sun UltraSPARC VI+
Processor Arch	64-bit	32/64-bit	32/64-bit	64-bit	64-bit	Dual 64-bit
Clock Rate	1.66GHz	3.66GHz	3.8GHz	700MHz	2.16GHz	1.5GHz
Cache (I/D/L2/L3)	16K/16K/ 256K/9M	12K/8K/ 1M/1M	12K/512K/ 2M	32K/32K	128K/128K/4M	64K/64K/ 2MB/32MB
Issue Rate/Core	6 issue	3 ROPs	3 ROPs	4 issue	8 issue	8 issue
Pipeline Stages	8 stages	22/24 stages	22/24 stages	6 stages	9 stages (int)	14 stages
Out of Order	None	126 ROPs	126 ROPs	48 instr	112 instr	None
Rename Regs	328 total	128 total	128 total	32/32	32/32	None
BHT Entries	512 x 2-bit	4K x 2-bit	4K x 2-bit	2K x 2-bit	16K x 2-bit	2 x 16 x 2-bit
TLB Entries	32L1I/32L1D/ 128L2I/128L2D	128I/64D	128I/64D	64 unified	(2,048+32)I/ (2,048+32)D	2 x (512+16)I/ 2 x (1,024+16)D
Memory B/W	10.6GB/s	5.3GB/s	6.4GB/s	1.6GB/s	4.3GBs	4.8GB/s
Package	mPGA-700	mPGA-604	mPGA-604	FCBGA-1153	LGA-908	FC-LGA 1368
IC Process	0.13µm 6M	0.09µm 6M	0.09µm 6M	0.11µm 7M	0.09µm 10M	0.09µm 9M
Die Size	432mm²	130mm²*	145mm²*	110mm²	294mm²	336mm²**
Transistors	592 million	125 million*	175 million*	7.2 million	400 million	295 million**
Est Die Cost	$165	$22	$24	$60	N/A	$125**
Power (Max)	130W	140W(TDP)	130W(TDP)	17W	65W	90W*
Availability	3Q05	2Q05	4Q05	1Q03	2Q05	3Q05
Configuration	1–256 way	1–8 way	1–2 way	1–512 way	1–128 way	4–72 way
SPEC_int2000(base)	1,490	1,388	1,810	N/A	1,456	N/A
SPEC_fp2000(base)	2,801	1,314	1,909	N/A	1,808	N/A

Source: vendors, except *In-Stat estimates. Estimated manufacturing cost does not include external cache chips. ** Contains two processors on one die. n/a = not available.

= IA-32

= IA-64

Current IA-32 Processors Manufacturing Cost

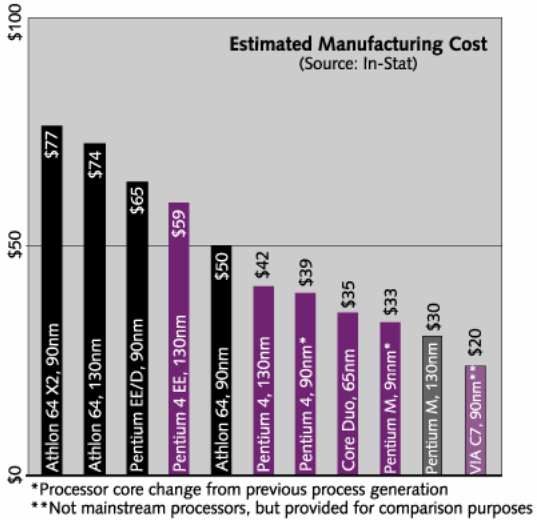

Source: Microprocessor Report, © Cahners 2006

Current IA-32 Processors Prices

Intel Processor	Cores/Frequency/L2/FSB/TDP	Process	Price	AMD Processor	Cores/Frequency/L2 Cache/HT/ MCT/TDP*	Process (SOI)	Price
Desktop PC							
Pentium EE 965	2/3.73GHz/2x2MB/1.06GHz/130W	65nm	$999	Athlon 64 FX-62	2/2.8GHz/2x1M/2.0GHz/667MHz/125W	90nm	$1,031
Pentium EE 955	2/3.46GHz/2x2MB/1.06GHz/130W	65nm	$999	Athlon 64 FX-60	2/2.6GHz/2x1M/2.0GHz/667MHz/89W	90nm	$827
Pentium D960	2/3.6GHz/2x2MB/800MHz/130W	65nm	$523	Athlon 64 X2 5000+	2/2.6GHz/2x1MB/2.0GHz/667MHz/89W	90nm	$696
Pentium D950	2/3.4GHz/2x2MB/800MHz/130W	65nm	$316	Athlon 64 X2 4800+	2/2.4GHz/2x1MB/2.0GHz/667MHz/89W	90nm	$645
Pentium D940	2/3.2GHz/2x2MB/800MHz/130W	65nm	$241	Athlon 64 X2 4600+	2/2.4GHz/2x512kB/2.0GHz/667MHz/89W	90nm	$558
Pentium D930	2/3.0GHz/2x2MB/800MHz/95W	65nm	$209	Athlon 64 X2 4400+	2/2.2GHz/2x1MB/2.0GHz/667MHz/89W	90nm	$470
Pentium 4 672	1/3.8GHz/2MB/800MHz/115W	90nm	$605	Athlon 64 4000+	1/2.4GHz/1MB/1.6GHz/400MHz/89W	90nm	$343
Pentium 4 662	1/3.6GHz/2MB/800MHz/115W	90nm	$401	Athlon 64 3800+	1/2.4GHz/512kB/2.0GHz/667MHz/89W	90nm	$290
Pentium 4 651	1/3.4GHz/2MB/800MHz/84W	90nm	$273	Athlon 64 3700+	1/2.4GHz/1MB/1.6GHz/400MHz/89W	90nm	$240
Pentium 4 641	1/3.2GHz/2MB/800MHz/84W	90nm	$218	Athlon 64 3500+	1/2.2GHz/512kB/2.0GHz/667MHz/89W	90nm	$189
Celeron D 355	1/3.33GHz/256kB/533MHz/73W	90nm	$79	Sempron 3600+	1/2.0GHz/256kB/1.6GHz/667MHz/62W	90nm	$123
Celeron D 351	1/3.2GHz/256kB/533MHz/73W	90nm	$69	Sempron 3500+	1/2.0GHz/128kB/1.6GHz/667MHz/62W	90nm	$109
Celeron D 346	1/3.06GHz/256kB/533MHz/73W	90nm	$59	Sempron 3400+	1/2.0GHz/256kB/1.6GHz/667MHz/62W	90nm	$97
Celeron D 341	1/2.93GHz/256kB/533MHz/73W	90nm	$59	Sempron 3200+	1/2.0GHz/128kB/1.6GHz/667MHz/62W	90nm	$87
Mobile PC							
Core Duo T2600	2/2.16GHz/2x1M/667MHz/31W	65nm	$637	Turion 64 X2 TL-60	2/2.0GHz/2x512kB/1.6GHz/667MHz/35W	90nm	$354
Core Duo T2500	2/2.0GHz/2x1M/667MHz/31W	65nm	$638	Turion 64 X2 TL-56	2/1.8GHz/2x512kB/1.6GHz/667MHz/35W	90nm	$263
Core Duo T2400	2/1.83GHz/2x1M/667MHz/31W	65nm	$639	Turion 64 X2 TL-52	2/1.6GHz/2x512kB/1.6GHz/667MHz/31W	90nm	$220
Core Duo T2300	2/1.66GHz/2x1M/667MHz/31W	65nm	$640	Turion 64 X2 TL-50	2/1.6GHz/2x256KB/1.6GHz/667MHz/31W	90nm	$184
Pentium M 780	1/2.26GHz/2MB/533MHz/27W	90nm	$637	Turion 64 MT-40	1/2.2GHz/1MB/1.6GHz/400MHz/25W	90nm	$225
Pentium M 770	1/2.13GHz/2MB/533MHz/27W	90nm	$423	Turion 64 MT-37	1/2.0GHz/1MB/1.6GHz/400MHz/25W	90nm	$189
Pentium M 760	1/2.0GHz/2MB/533MHz/27W	90nm	$294	Turion 64 MT-34	1/1.8GHz/1MB/1.6GHz/400MHz/25W	90nm	$159
Pentium M 750	1/1.80GHz/2MB/533MHz/27W	90nm	$241	Turion 64 MT-32	1/1.8GHz/512KB/1.6GHz/400MHz/25W	90nm	$150
Celeron M 380	1/1.6GHz/1MB/400MHz/21W	90nm	$134	Mobile S3300+	1/2.0GHz/128kB/1.6GHz/400MHz/25W	90nm	$113
Celeron M 370	1/1.5GHz/1MB/400MHz/21W	90nm	$107	Mobile S3100+	1/1.8GHz/256kB/1.6GHz/400MHz/25W	90nm	$98
Celeron M 360	1/1.4GHz/1MB/400MHz/21W	90nm	$86	Mobile S3000+	1/1.8GHz/128kB/1.6GHz/400MHz/25W	90nm	$88
Celeron M 350	1/1.3GHz/1MB/400MHz/21W	90nm	$86	Mobile S2800+	1/1.6GHz/256kB/1.6GHz/400MHz/25W	90nm	$77

*The AMD processors feature an integrated north bridge. As a result, the I/O bus (HT for HyperTransport) and memory bus (MCT) are listed instead of the front side bus (FSB). SXXXX = Sempron.

Source: Microprocessor Report, © Cahners 2006

Outline

- Historical limitations of IA-32
- How some Pentium designs have worked around the main limitations
  - PentiumPro: Achieving superscalar out-of-order execution on a CISC
  - Pentium4: Achieving 2GHz clock frequency

Legacy IA-32 Features

- Very small number of registers, partly dedicated or specialised
- Natively 16-bit, extended to 32 in successive steps requiring backward compatibility (e.g., 3 modes for address generation)
- Highly variable instruction length and encoding (1 to 17 bytes in original IA-32, prefixes, postfixes, etc.)
- CISC instruction set

Registers (I)

- Very small number of general purpose registers (approx. 4 integer plus 8 FP—not shown, versus 32+32 typ. RISC)

Segment Registers

	15 … 0	
CS		Code segment ptr
SS		Stack segment ptr
DS		Data segment ptr
ES		Extra data segm. ptr
FS		Data segment ptr 2
GS		Data segment ptr 3

General Registers + PC + Flags

31 … 15	15 … 8	7 … 0	
EAX	AX	AH / AL	Accumulator
ECX	CX	CH / CL	Count reg: string, loop
EDX	DX	DH / DL	Data reg: multiply, divide
EBX	BX	BH / BL	Base addr reg
ESP	SP		Stack pointer
EBP	BP		Base ptr (base stack reg)
ESI	SI		Index reg, string src ptr
EDI	DI		Index reg, string dest ptr
EIP	IP		Instruction ptr (**PC**)
EFLAGS	FLAGS		Condition codes

◆ → **80386** ◆ → **8086**

Registers (II)

- Small number of registers makes spilling more frequent
- Advanced compiler techniques (e.g., loop unrolling, Lesson 10) increase register pressure
- Partial specialization of the registers makes effective compiler scheduling difficult

Memory Addressing (I)

- Real Mode (8086)

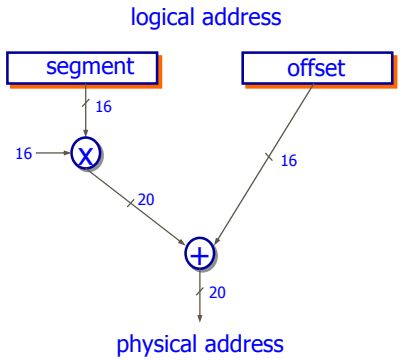

Memory Addressing (II)

- Protected Mode (80286)

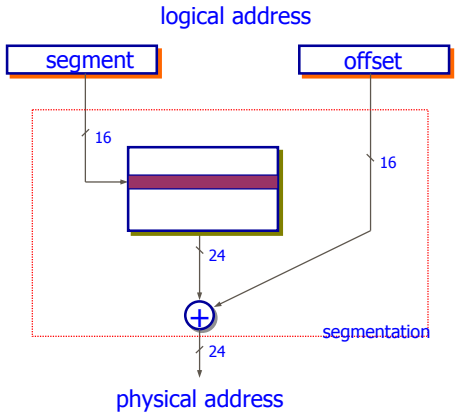

Memory Addressing (III)

- Protected Mode (80386, 80486, and Pentium)

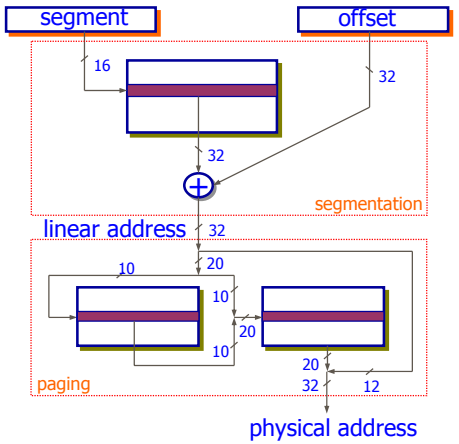

Addressing Modes (I)

- **Absolute**
- **Register indirect** → [reg]
  - 16-bit registers: BX, SI, DI
  - 32-bit registers: EAX, ECX, EDX, EBX, ESI, EDI
- **Displacement** → [reg + *displacement*]
  - 16-bit registers: BP, BX, SI, DI
  - 32-bit registers: EAX, ECX, EDX, EBX, ESI, EDI
  - Displacement on 8, 16, or 32 bits
- **Indexed** → [base reg + reg]
  - 16-bit registers: BX+SI, BX+DI, BP+SI, BP+DI

Addressing Modes (II)

- **Indexed with displacement**
  → [base reg + reg + *displacement*]
  - Same registers as in mode indexed
- **Scaled indexed** → [base reg + 2^{scale} x reg]
  - Only in 32-bit mode
  - Scale is 0, 1, 2, or 3
  - Index register can be any of the basic registers (except ESP)
  - Base register can be any of the basic registers
- **Scaled indexed with displacement**
  → [base reg + 2^{scale} x reg + *displacement*]

Address Segment

- For every indirect addressing (e.g., [reg]) the appropriate segment would be needed
- Default:
  - References to instructions (IP) use CS (code segment register)
  - References to stack (BP or SP) use SS (stack segment register)
  - All other references use DS (data segment register)
- A one-byte instruction prefix can modify the default

Instructions—IA-32 is not the same architecture since the mid-80's

- Classic CISC set derived from extended accumulator architecture
- Improved orthogonality in the 32-bit extensions (80386)
- Added FP capabilities previously on a coprocessor (80486)
- Added MultiMedia Extensions **MMX** as SIMD (single-instruction multiple-data) integer instructions (Pentium II)
- Added Streaming SIMD Extension **SSE**, most notably consisting of SIMD FP instructions (Pentium III)
- Added **SSE2**, essentially extension of MMX+SSE to 128 bits (Pentium 4)

Operand Types

- **Not** a Load/Store architecture

Source 1 = Destination	Source 2
Register	Register
Register	Immediate
Register	Memory
Memory	Register
Memory	Immediate

- Immediate values can be on 8, 16, or 32 bits

Instruction Examples

JE addr	if equal(CC) then IP ← addr (IP-128 ≤ addr < IP+128)
JMP addr	IP ← addr
CALL addr,seg	SP ← SP-2; Mem[SS:SP] ← IP+5 SP ← SP-2; Mem[SS:SP] ← CS IP ← addr; CS ← seg
MOVW BX,[DI+45]	BX ← Mem[DS:DI+45]
PUSH SI	SP ← SP-2 Mem[SS:SP] ← SI
POP DI	DI ← Mem[SS:SP] SP ← SP+2
ADD AX,#6765	AX ← AX+6765
TEST DX,#42	set CC flags with (DX and 42)
MOVSB	Mem[ES:DI] ← Mem[DS:SI] DI ← DI+1 SI ← SI+1

Instruction Encoding

- One instruction coded on 1 to 17 bytes in original IA-32
- Several types of modifiers/prefixes
- Two combinations of constants of variable length
  - Immediate and Displacement
  - 8, 16, and 32-bit
- Opcode "lost" and only moderately orthogonal

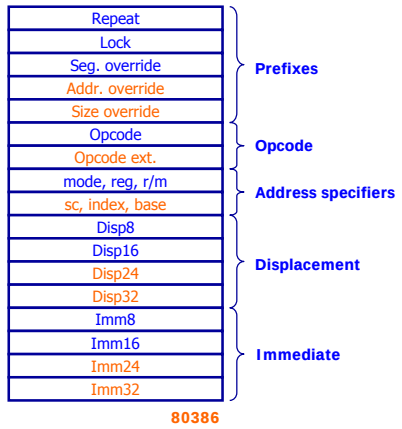

Examples of Instruction Encoding

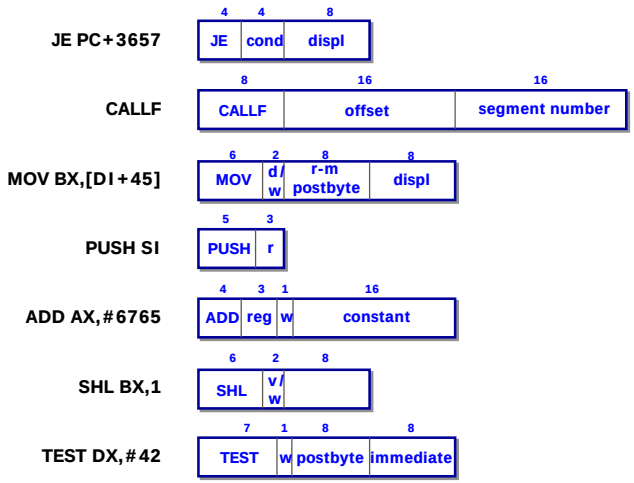

1995: PentiumPro (P6) A Superscalar IA-32 CISC?

- How to adapt the superscalar ideas to fit such an irregular architecture?
  - Complexity of decoding is huge
  - Parallel decoding of instructions is tough due to an encoding strongly variable in lenght
  - Instructions mix memory operations with computations
  - Too few registers

PentiumPro Microarchitecture: Out-of-order CISC Execution

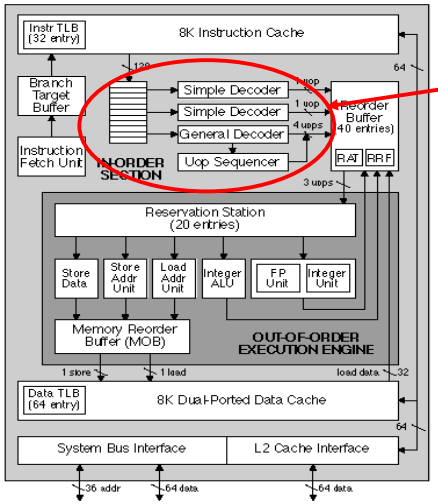

Decoder from IA-32 to P6 uops: A "CISC-to-RISC converter"...

Classic RISC out-of-order engine

Source: Microprocessor Report, © Cahners 2001

PentiumPro In-Order Section

- Converts every IA-32 instruction into one or more internal RISC-like 118-bit instructions (micro-operations or uops); on average 1 instruction = 1.5-2.0 uops
- Three decoders and a sequencer work in parallel to perform the conversion
  - Two highest priority simple decoders intercept register-register operations (1 instr. → 1 uop)
  - A low priority general decoder handles all other basic operations (1 instr. → 4 uops)
  - A sequencer is used by the general decoder for very complex operations (1 instr → several groups of 4 uops)
- Reorder Buffer (ROB) implements renaming and commits uops in program order to the Real Register File (RRF)

PentiumPro Out-of-order Section

- Superscalar very similar to the general model studied in previous lessons
- Up to 20 uops wait in the Reservation Stations until the operands are all available
- A maximum of 5 uops can be issued per cycle: a generic calculation (int or FP), a simple integer (no shift, mul, nor div), a load, a store address, and a store data
- A Memory Reorder Buffer (MOB) reorders memory accesses and waits for D-cache availability

PentiumPro Die
300mm² 0.5µm 4ML BiCMOS

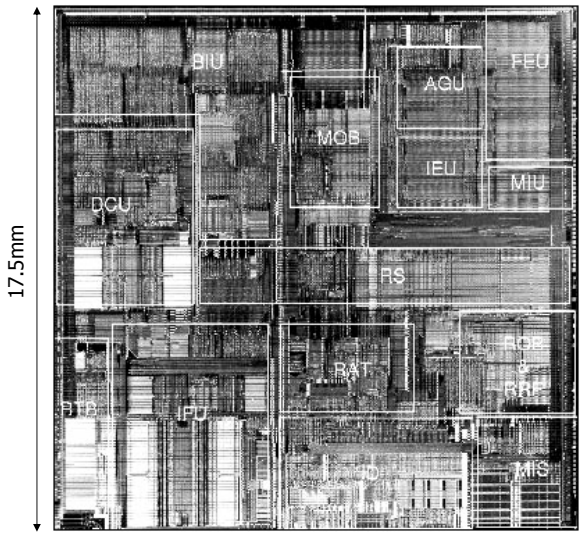

Source: Microprocessor Report, © Cahners 2001

PentiumPro Package

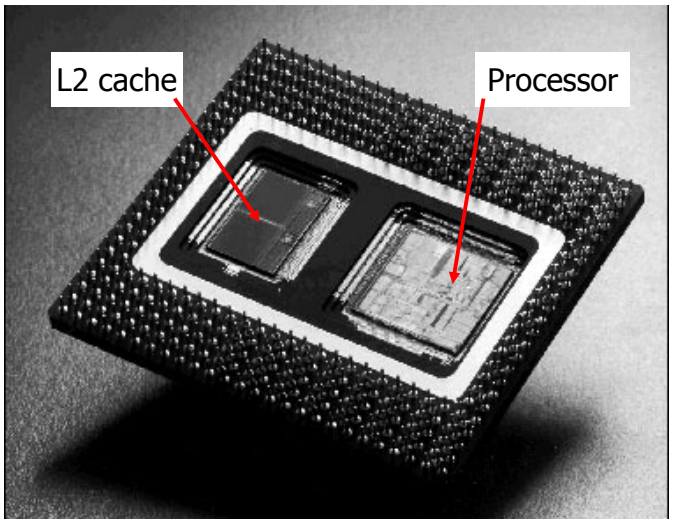

Source: Microprocessor Report, © Cahners 2001

2000: Pentium 4
Processors for the Multi-GHz Era

- Isn't Pentium dead in favour of IA-64 and Itanium? Clearly not…
- How to modify Pentium III to achieve way less than 1ns of cycle time?
  - Pipeline expansion? (Pentium III: 10 stages)
  - Wire propagation time becomes very tangible compared to computation
  - uop decoding very heavy

Evolution of Pentium Pipeline: From 5 to 20 Stages

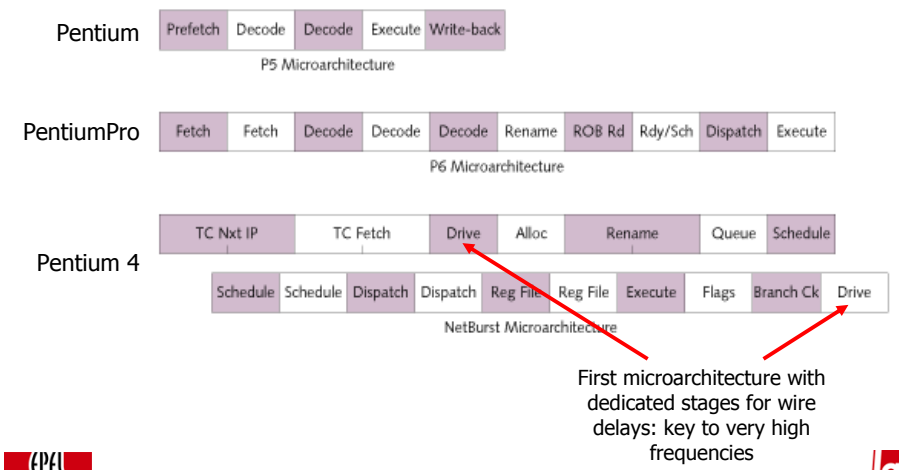

Source: Microprocessor Report, © Cahners 2001

Some Pentium 4 Characteristics

- **Trace caches** to memorize approx. 12,000 recent uops—sort of L0 cache to avoid IA-32 instruction decoding from the main loop
- Approx. 126 uops can be in-flight at one time (3 times more than Pentium III)
- **Data speculation** to execute a potentially dependent load before a store: if the dependence was real, the load is squashed and replayed
- P4 ALUs can perform simple operations in half clock-cycle, to sustain throughput
  - Two dependent operations can be scheduled in the same cycle

AMD Hammer x86-64: Extension of IA-32 to 64 bits (I)

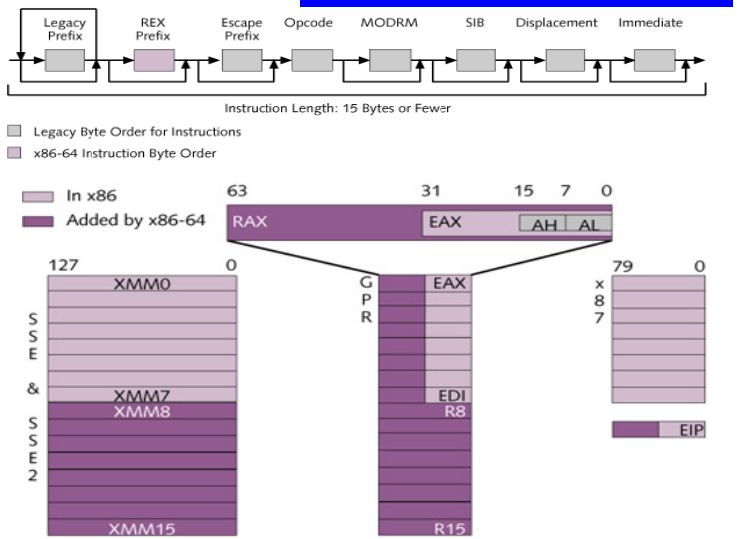

Source: Microprocessor Report, © Cahners 2000-01

AMD Hammer x86-64: Extension of IA-32 to 64 bits (II)

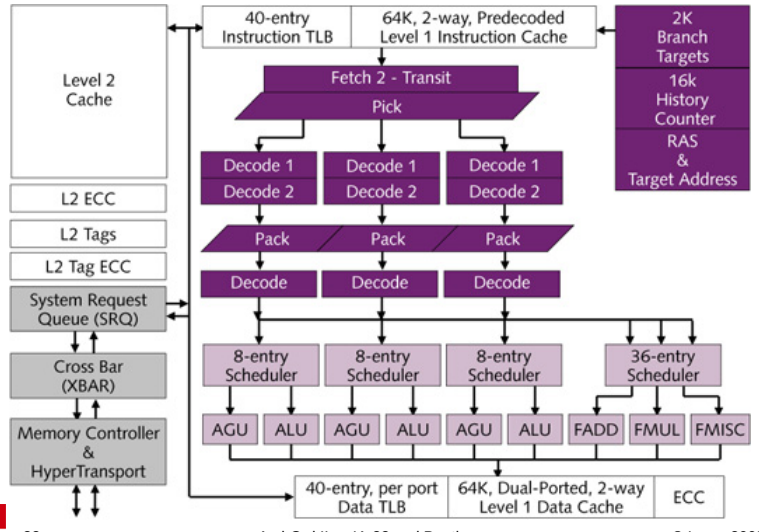

Source: Microprocessor Report, © Cahners 2001

Conclusions

- IA-32 is the oldest important ISA around
- It is not absolutely fixed but constantly evolving with many new add-ons (MMX, SSE, SSE2, CMOV, etc.)
- Intel has managed to continue pushing the performance by adapting to its CISC nature the techniques developed to speed-up newer RISC processors
  - Similar work has been done by some competitors—notably by AMD with Athlon, for a few months the fastest IA-32 processor on the market
- Is IA-32 really dying? Is it evolving toward 64 bits?!…

References and Where to Learn More

- References:
  - COD, Sections 3.12, 4.9, 5.7, 6.9, and 7.6
- Where to learn more:
  - Stallings, Computer Organization & Architecture, 5th ed., 2000, Section 13.3
  - D. Alpert and D. Avnon, *Architecture of the Pentium Microprocessor*, IEEE Micro, June 1993
  - L. Gwennap, *Intel's P6 Uses Decoupled Superscalar Design*, Microprocessor Report, 16th February 1995
  - P. Glaskowsky, *Pentium 4 (Partially) Previewed*, Microprocessor Report, 28th August 2000
  - S. Leibson, AMD Drops 64-bit Hammer on x86, Microprocessor Report, 4th September 2000

All papers available at http://lap.epfl.ch/courses/archord2/