

The Intel P6 and NetBurst microarchitectures

M. Sonza Reorda

Politecnico di Torino
Dipartimento di Automatica e Informatica

1

P6 microarchitecture

The first Pentium Pro (based on the P6 microarchitecture) was released in 1995.

It was the first real superscalar processor in the x86 family. Its predecessor (P5) included a double pipeline.

To allow both adopting a superscalar architecture and maintaining software compatibility, translation of IA-32 instructions into simple RISC instructions (called μ ops) is introduced.

2

P6

The Intel P6 microarchitecture is the basis for

- PentiumPro
- Pentium II
- Pentium III.

The 3 processors differ because of

- Some instruction extensions (MMX, SSE)
- Clock rate
- Cache architecture
- Memory interface.

3

Evolution in P6 processors

Processor	First ship date	Clock rate range	L1 cache	L2 cache
Pentium Pro	1995	100–200 MHz	8 KB instr. + 8 KB data	256 KB–1024 KB
Pentium II	1998	233–450 MHz	16 KB instr. + 16 KB data	256 KB–512 KB
Pentium II Xeon	1999	400–450 MHz	16 KB instr. + 16 KB data	512 KB–2 MB
Celeron	1999	500–900 MHz	16 KB instr. + 16 KB data	128 KB
Pentium III	1999	450–1100 MHz	16 KB instr. + 16 KB data	256 KB–512 KB
Pentium III Xeon	2000	700–900 MHz	16 KB instr. + 16 KB data	1 MB–2 MB

4

Ev

The CPU (5.5 million transistors) and the L2 cache (2 million transistors) are integrated into a single package (*cartridge*) composed of two dies.

The two parts are connected by a dedicated bus running at the full CPU speed (200 MHz).

Processor	First ship date	Clock rate range	L1 cache	L2 cache
Pentium Pro	1995	100–200 MHz	8 KB instr. + 8 KB data	256 KB–1024 KB
Pentium II	1998	233–450 MHz	16 KB instr. + 16 KB data	256 KB–512 KB
Pentium II Xeon	1999	400–450 MHz	16 KB instr. + 16 KB data	512 KB–2 MB
Celeron	1999	500–900 MHz	16 KB instr. + 16 KB data	128 KB
Pentium III	1999	450–1100 MHz	16 KB instr. + 16 KB data	256 KB–512 KB
Pentium III Xeon	2000	700–900 MHz	16 KB instr. + 16 KB data	1 MB–2 MB

5

Ev

MMX extension

6 processors

Processor	First ship date	Clock rate range	L1 cache	L2 cache
Pentium Pro	1995	100–200 MHz	8 KB instr. + 8 KB data	256 KB–1024 KB
Pentium II	1998	233–450 MHz	16 KB instr. + 16 KB data	256 KB–512 KB
Pentium II Xeon	1999	400–450 MHz	16 KB instr. + 16 KB data	512 KB–2 MB
Celeron	1999	500–900 MHz	16 KB instr. + 16 KB data	128 KB
Pentium III	1999	450–1100 MHz	16 KB instr. + 16 KB data	256 KB–512 KB
Pentium III Xeon	2000	700–900 MHz	16 KB instr. + 16 KB data	1 MB–2 MB

6

Evolution of SSE extension in Pentium processors

Processor	Year	Clock rate range	L1 cache	L2 cache
Pentium Pro	1995	100–200 MHz	8 KB instr. + 8 KB data	256 KB–1024 KB
Pentium II	1998	233–450 MHz	16 KB instr. + 16 KB data	256 KB–512 KB
Pentium II Xeon	1999	400–450 MHz	16 KB instr. + 16 KB data	512 KB–2 MB
Celeron	1999	500–900 MHz	16 KB instr. + 16 KB data	128 KB
Pentium III	1999	450–1100 MHz	16 KB instr. + 16 KB data	256 KB–512 KB
Pentium III Xeon	2000	700–900 MHz	16 KB instr. + 16 KB data	1 MB–2 MB

7

μ ops

The P6 translates each IA-32 instruction into a series of micro-operations (μ ops), which are similar to typical RISC instructions.

Up to 3 IA-32 instructions are fetched, decoded, and translated into μ ops at each clock cycle.

8

μops execution

μops are executed by an out-of-order speculative pipeline using register renaming and ROB.

Up to 3 μops per clock can be renamed and dispatched to the reservation stations.

Up to 3 μops per clock can be committed.

9

Pipeline

It is composed of 14 stages:

- 8 stages are used for in-order instruction fetch, decode, and dispatch**
- 3 stages are used for out-of-order execution in one out of 5 functional units**
- 3 stages are used for instruction commit.**

A stage stalls if the input buffer does not provide it with new data, or if the output buffer is full.

10

Fetch unit

It fetches one cache line (32 bytes) per clock cycle from the L1 code cache.

It identifies instruction boundaries, and performs branch prediction.

11

Decoding unit

It is composed of 3 units working in parallel: two for simple instructions and one for complex instructions.

Simple instructions are mapped to single μ -ops. Complex instructions are mapped to several μ -ops (up to 4): this procedure requires more than one clock cycle.

The decoding stage can thus produce up to 6 μ -ops per clock cycle, which are then written to the *instruction pool*.

12

Instruction pool

It stores up to 20/30 μ -ops.

It continuously checks which μ -ops can be sent to the functional units.

μ -ops can be sent to functional units in any order.

13

Branch prediction

It is based on a Branch Target Buffer (BTB) with 512 entries.

If the BTB fails, a static prediction is exploited, i.e.,

- Backward branches are predicted taken
- Forward branches are predicted untaken.

The penalty for mispredicted branches is between 10 and 15 cycles, plus the overhead for incorrectly speculated instructions.

14

Register alias table

It converts references to x86 registers to references to the 40 internal registers, thus performing *register renaming*.

15

Dispatch unit

It assigns instructions

- to one of the 20 reservation stations *and*
- to one of 40 entries in the ROB.

16

Execution units

They are

- Two units for FP instructions (FPUI, FPUII)
- Two units for integer instructions (INTI, INTII)
- One memory interface unit (MIU).

The 5 units can operate in parallel and independently.

17

Latency and repeat rate

Instruction name	Pipeline stages	Repeat rate
Integer ALU	1	1
Integer load	3	1
Integer multiply	4	1
FP add	3	1
FP multiply	5	2
FP divide (64-bit)	32	32

18

Retire unit

It continuously looks in the instruction pool (playing the role of the ROB) for completed instructions.

Retirement happens according to a strict in-order strategy.

As soon as it becomes possible, it performs

- Data dependencies resolution
- Branch prediction verification
- Real register updating.

19

Performance measurements

The following data come from a paper by Bhandarkar and Ding (1997).

In the case of P6, the ideal CPI is 0.33.

The P6 can fall behind because at some clock cycles it failed in completing 3 instructions per clock cycle. This can be due to a number of problems occurred in any of the previous stages.

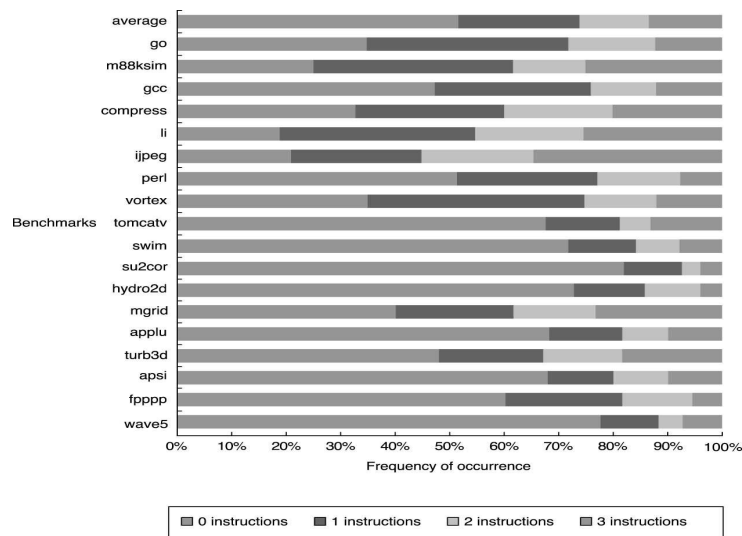
20

Possible causes for stalls

- Instruction cache miss
- The 3 decoded IA-32 instructions generated more than the maximum number of μ ops
- Lack of reservation stations or ROB entries
- Data dependency
- Data cache miss
- Branch mispredictions.

21


Stalls in the decode cycle



22

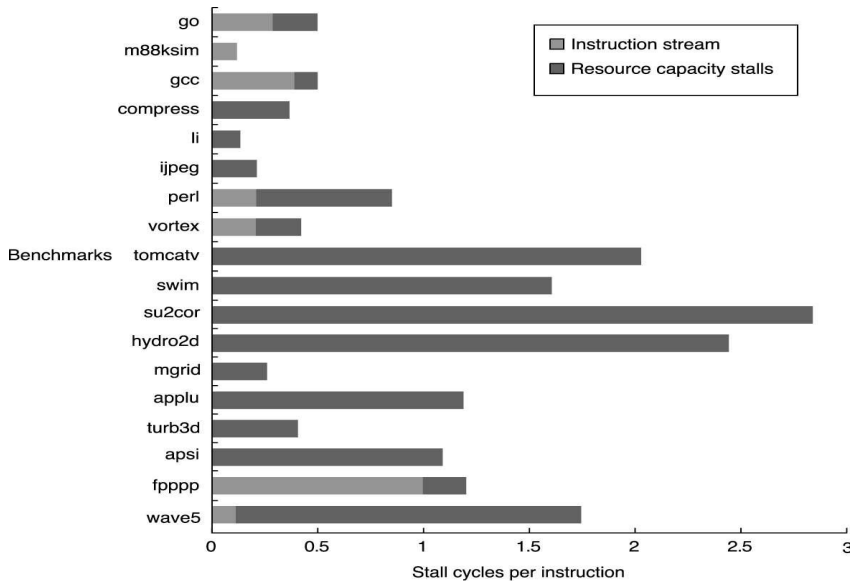
Stalls in the decode

In the average, 0.87 instructions are decoded per cycle




23

Causes of stalls in the decode cycle



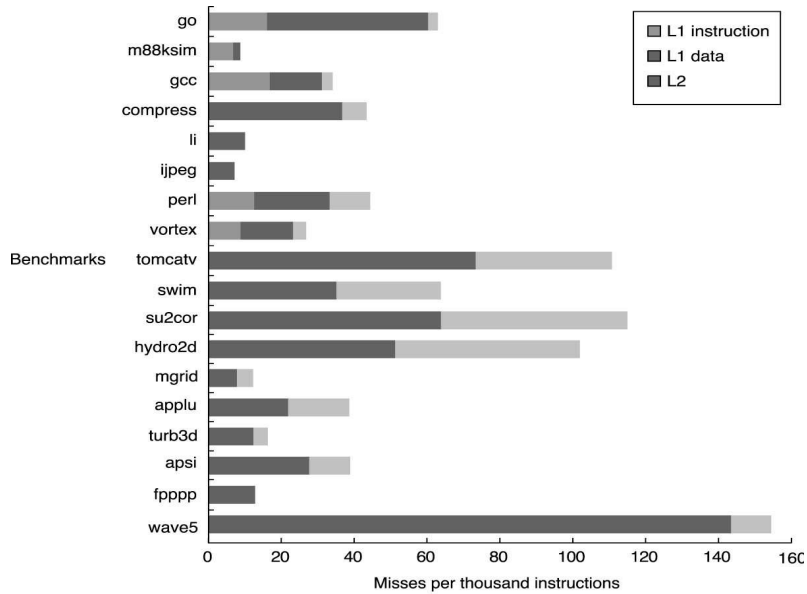
24

Number of μ ops per instruction



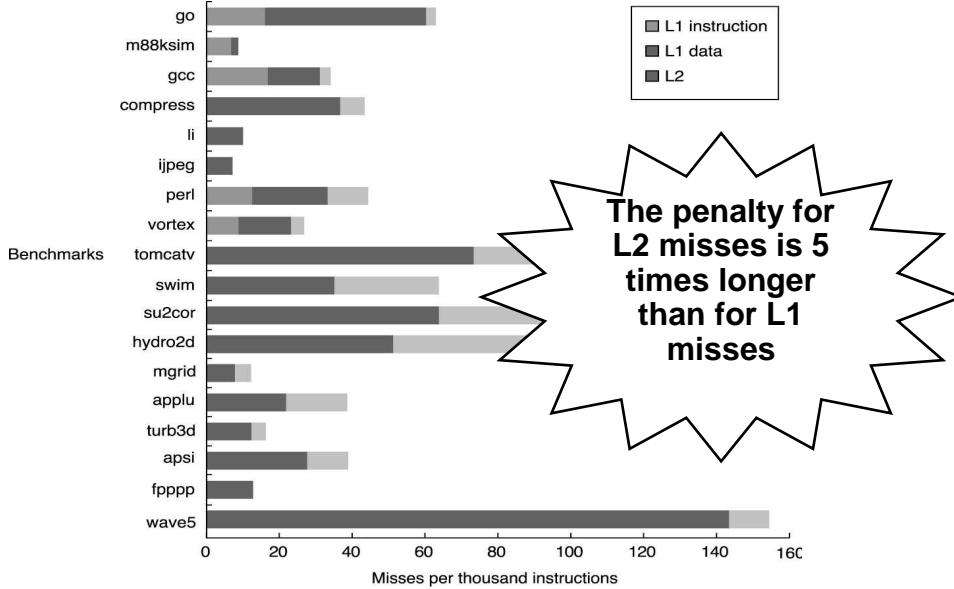
25

Cache miss ratio



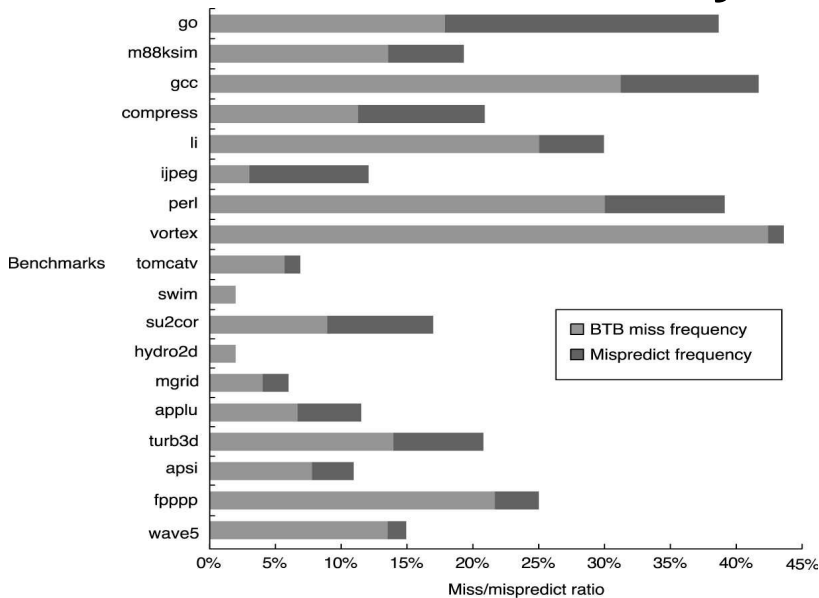
26

Cache miss ratio



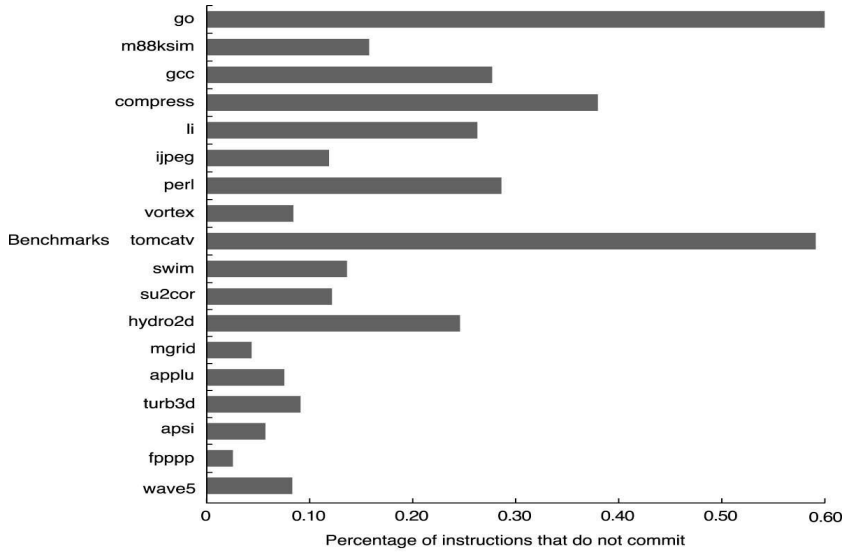
27

Prediction accuracy



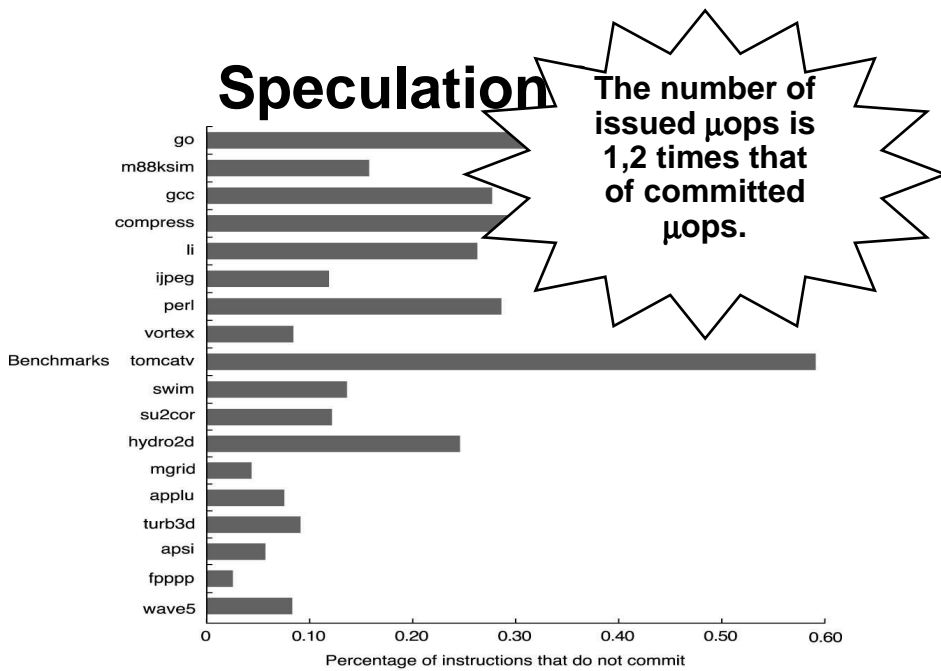
28

Speculation factor



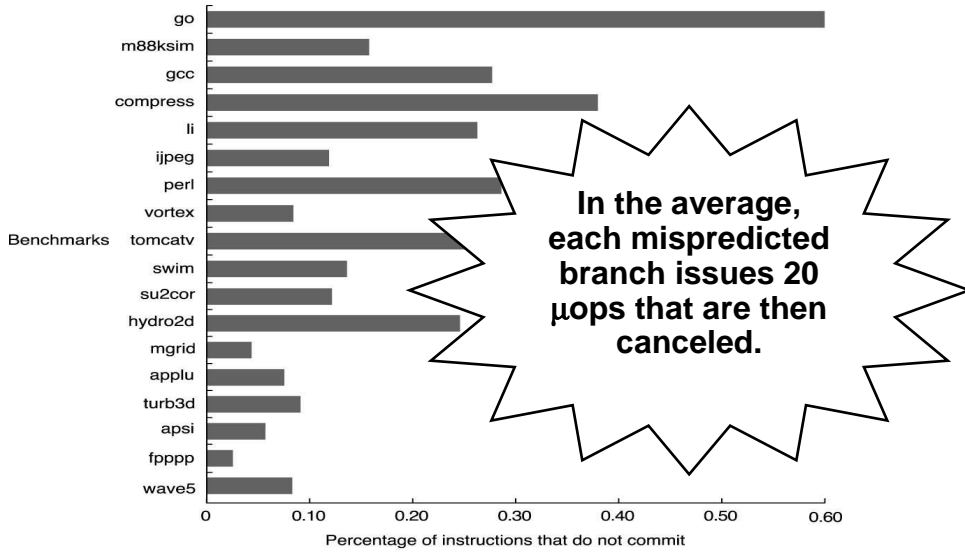
29

Speculation




30

Speculation factor



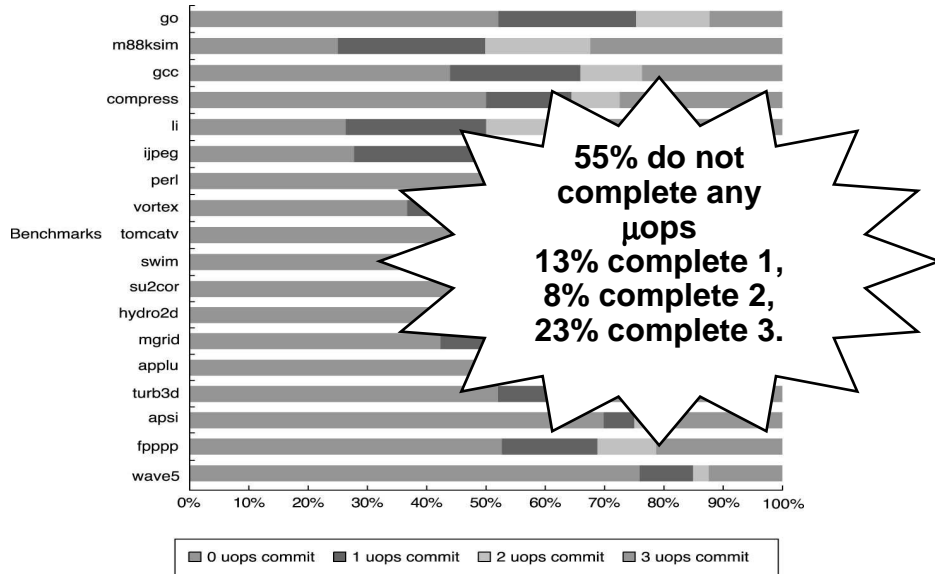
31

Overall performance



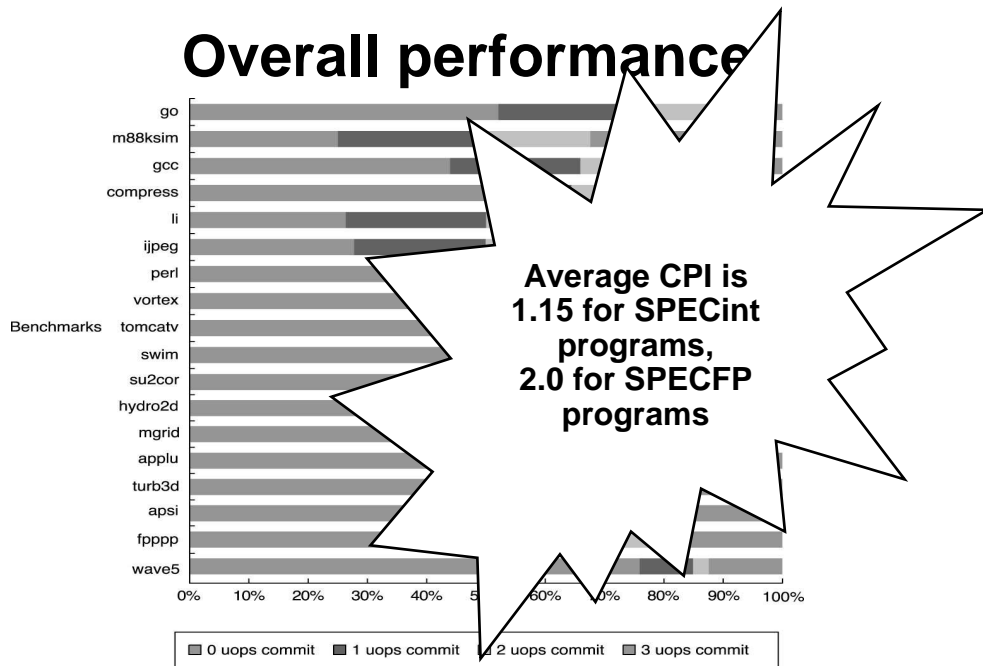
32

Overall performance



33

Overall performance



34

The NetBurst architecture

NetBurst is the name of the Pentium 4 microarchitecture, first introduced in 2000.

The two main goals of this architecture are

- To support a higher clock frequency
- To maintain the sustained execution throughput as close as possible to the ideal maximum.

NetBurst still maintains the idea of translating IA-32 instructions into μ ops.

However, to achieve the above goals, several differences were introduced in the NetBurst architecture with respect to the P6 one.

35

P6/NetBurst differences (I)

- Deeper pipeline: initial NetBurst architectures required 21 cycles for an add instruction, while P6 only required 10 cycles; this allowed increasing the operating frequency to 1.5 GHz. Further versions increased the pipeline depth to 31 stages, and the frequency to 3.2 GHz (2004).
- NetBurst uses a 128 entry register renaming rather than the reorder buffer with 40 entries used in P6
- NetBurst has 7 integer execution units instead of the 5 of P6 (one more integer ALU and one more address computation unit)

36

P6/NetBurst differences (II)

- *Trace cache* to improve instruction fetch performance
- More effective BTB: eight time larger and with an improved prediction algorithm; another BTB works on μ ops
- Improved IA-32 instruction decoding: when a single instruction corresponds to more than 3 μ ops, the μ op sequence is generated from a microcode ROM

37

P6/NetBurst differences (III)

- Faster ALU (0.5 cycles vs. 1) and faster data cache (2 cycles vs. 3)
- SSE2 floating point instructions, making the Pentium 4 more effective when executing floating point programs
- Pentium 4 supports multithreading.

38

Dispatch unit

At every clock cycle, up to 6 μ ops can be dispatched to the functional units by the 4 dispatch ports:

- One for the load unit
- One for the store unit
- One for basic ALU operations
- One for FP and integer operations.

39

Trace cache

A Trace Cache is a type of instruction cache that holds sequences of instructions (including sequences of nonadjacent instructions separated by branches).

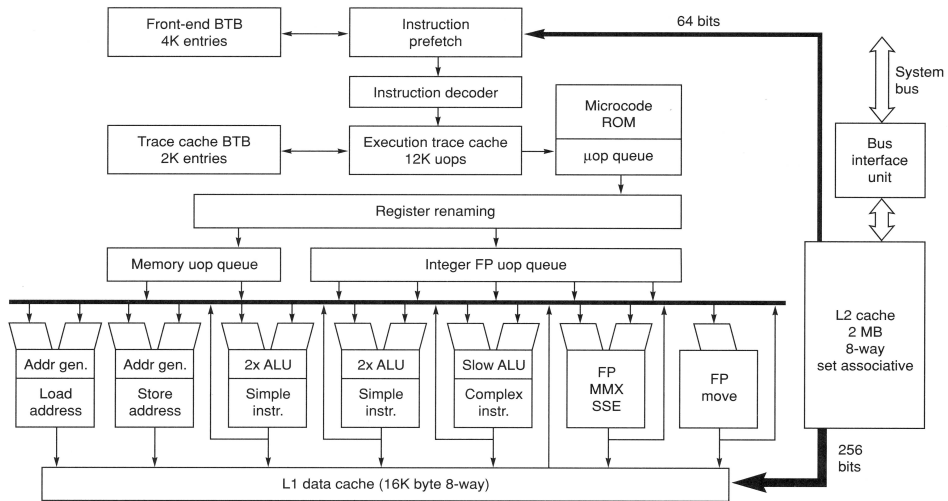
In the Pentium 4

- The trace cache contains μ ops
- Fetch is performed from the trace cache wherever possible
- Only on a trace cache miss IA-32 instructions are fetched from the L2 cache and decoded.

The hit rate of the trace cache is often higher than 99%.

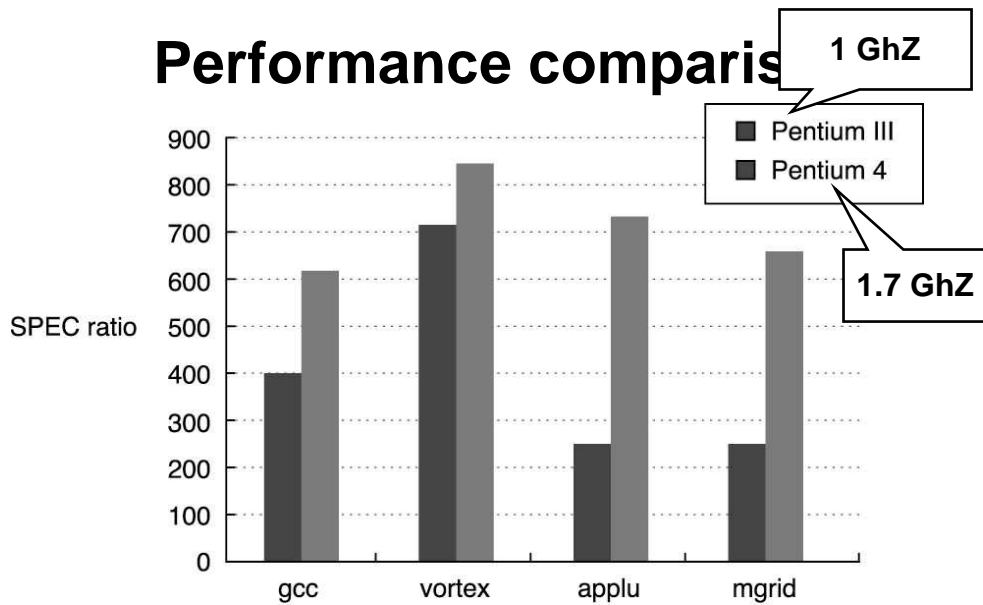
40

Architecture

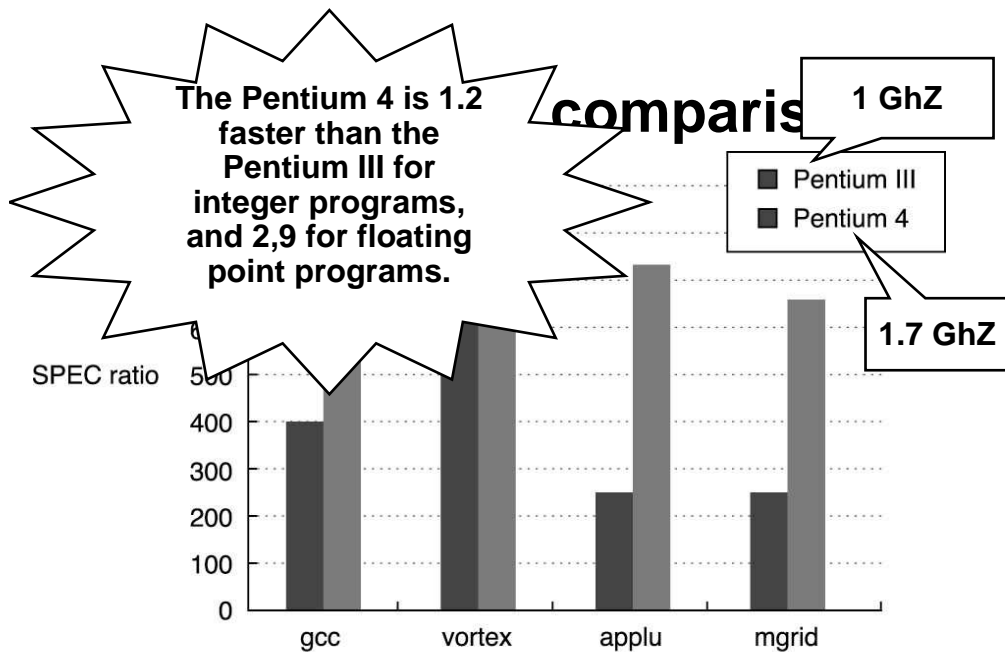


41

Performance comparison



42



43

IBM Power5

It was designed for high-performance integer and FP applications.


It contains two processor cores, each capable of sustaining four instructions per clock (two FP and two load-store instructions).

Its highest clock rate (2005) is 1.9 GHz.

For sake of comparison, Pentium 4 can sustain 3 instructions per clock, with a very deep pipeline a max clock of 3.8 GHz.

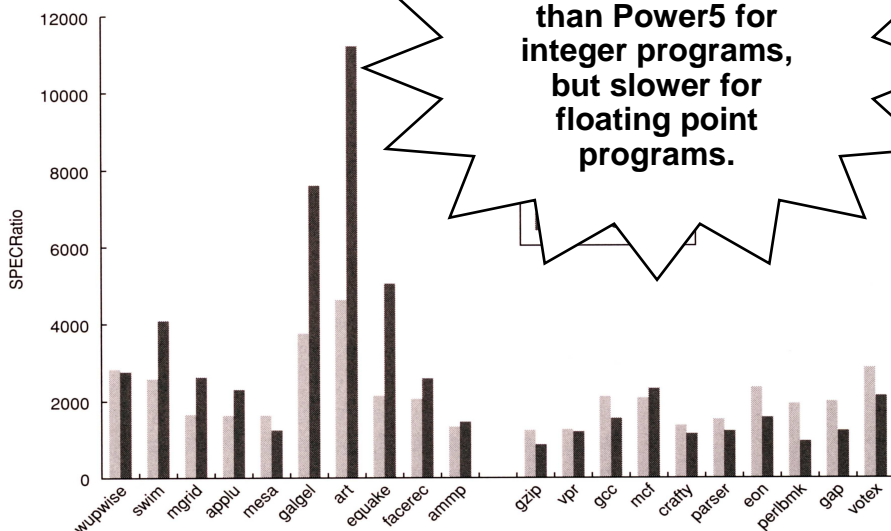
44

Performance comparison



45

Performance



46