IA-32

Nicholas FitzRoy-Dale

Datapoint 2200

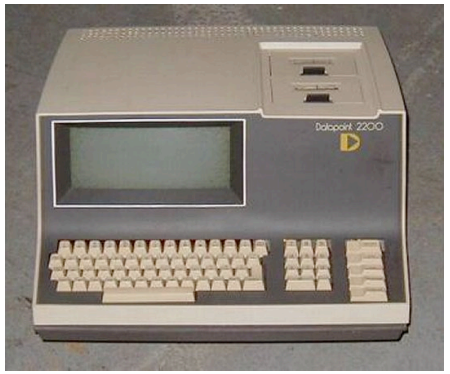

- Released 1970
- 2K shift register main memory
- CPU: ~100 TTL components
- Instruction set implemented by Intel in the 8008

History

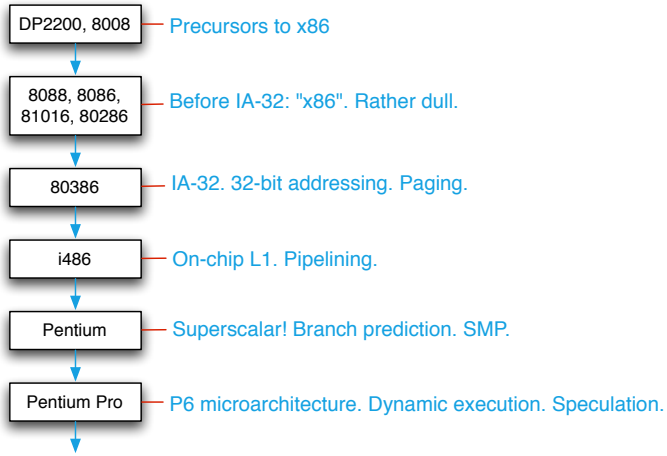

History

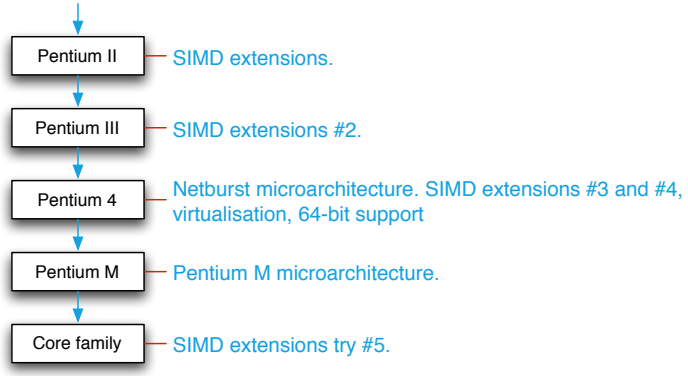

Legacy

- CISC ISA
- Lack of registers
- Four modes of operation
- Segmentation
- Superscalar
- Complex addressing modes

Microarchitectures

- New one “every two years”
- Focus on P6

P6 microarchitecture

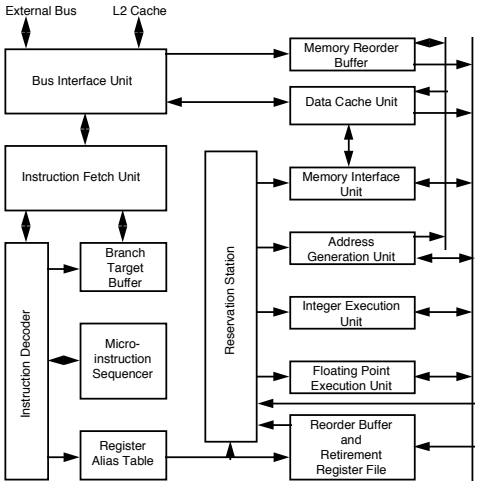

P6 microarchitecture

- As RISC as possible
- Register renaming
- Superscalar
- Out-of-order execution
- Speculation

P6: As RISC as possible

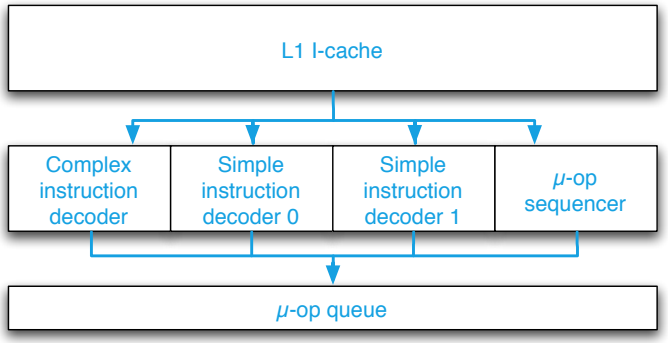

Instruction decode

- Pentium 4: Trace cache
- Pentium M, Core: μ-op fusion
- ARM: Single instruction per clock
- IA-64: Two identical decoders (currently)

P6: "Register renaming"

- Performed by ROB
- 40 physical registers in RAT
- ARM: Lots of GPRs
- IA-64: Lots of GPRs and register rotation

Avoiding stalls

- Out-of-order execution
- Branch prediction
- Speculation

Out-of-order execution

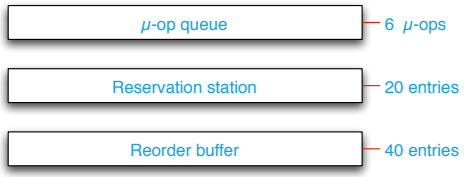

P6: Superscalar

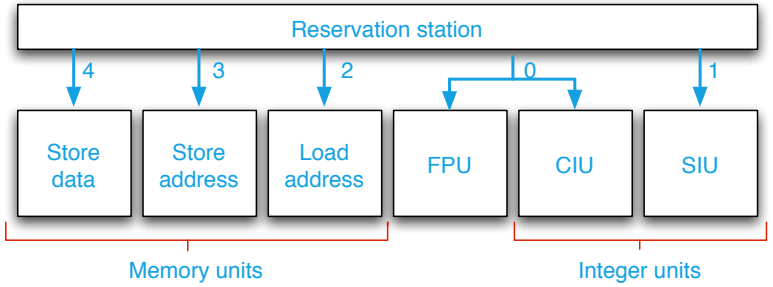

Superscalar: Netburst

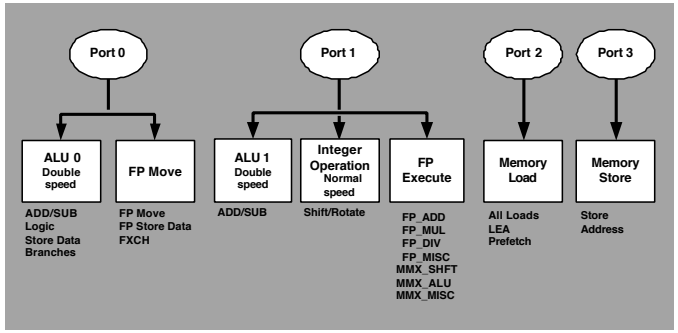

Superscalar

- ARM:
  - Varies per-core:
    - None
    - Separate ALU, MAC and LS (ARM11)
- IA-64:
  - 2 I-units, 2 F-units, 3 B-units
  - But very few templates avoid “split issue”

P6: Branch prediction

- Branch target buffer: 512 entries
  - Branch history and predicted address
  - Mispredicts: 10-15 cycles
- Static prediction
- ~90% hit rate
- Pentium M: Loop detection

P6: Speculation

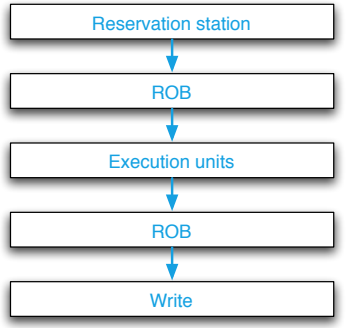

Registers

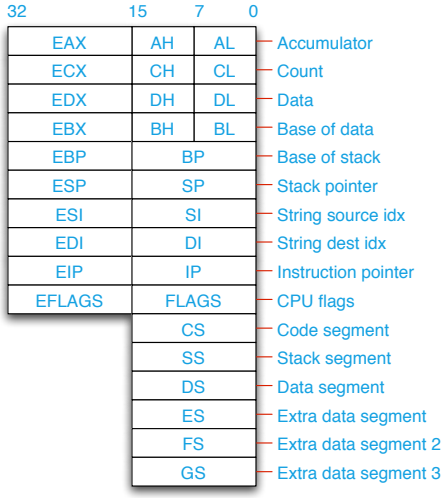

Floating-point

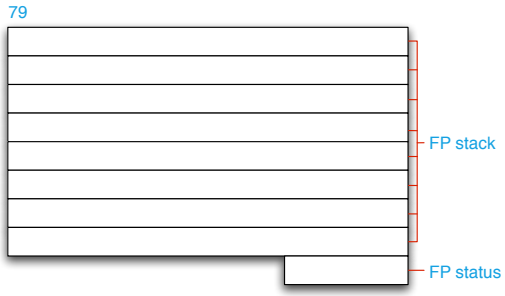

MMX

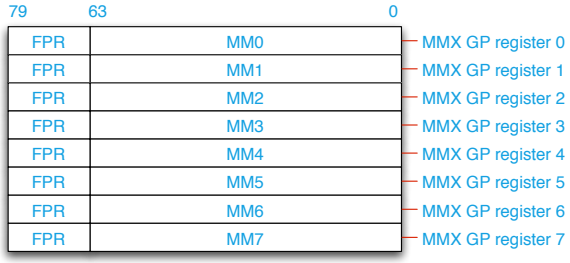

SSE, SSE2, SSE3, SSE4

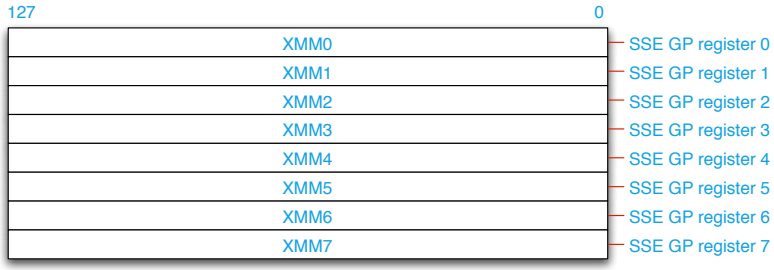

IA-32 Instructions

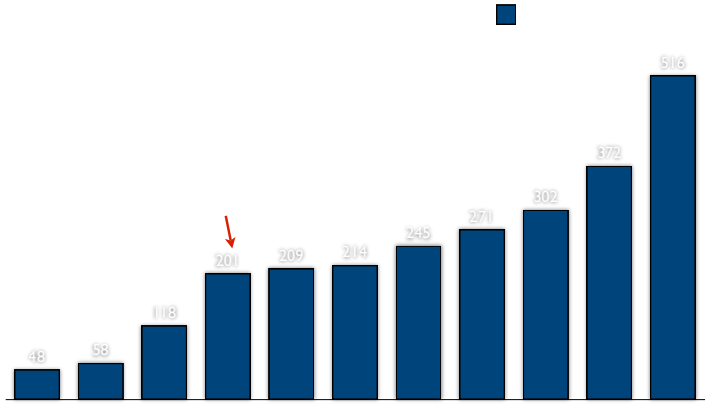

Instruction counts

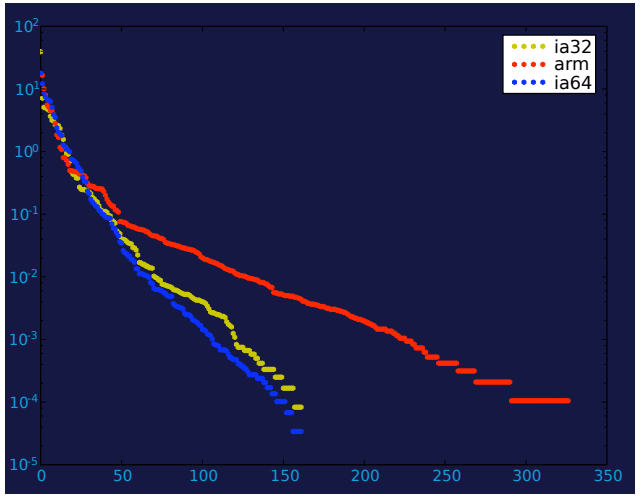

Instruction counts

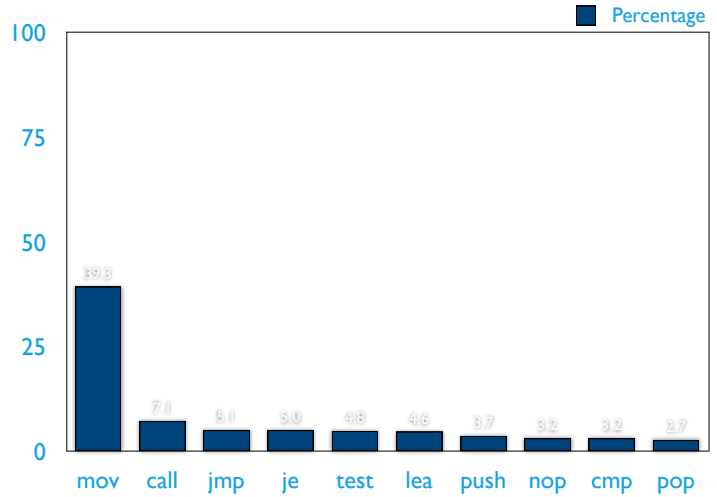

IA-32 vs ARM

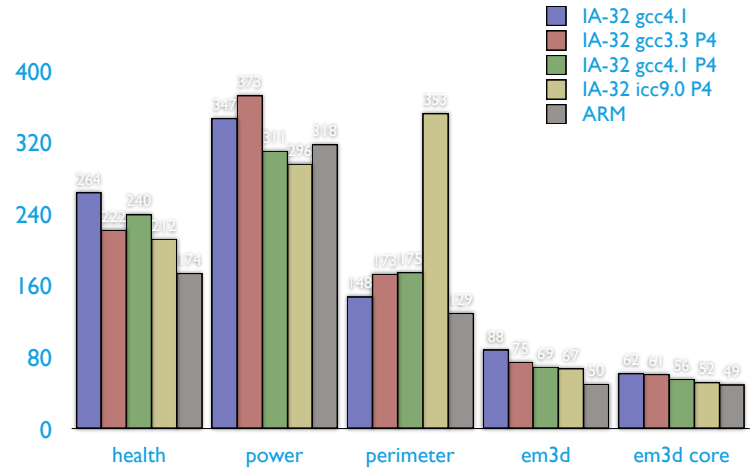

compute_nodes()

The code	ARM	IA-32
other_value = localnode->from_values[i];	ldr r4, [r0, #12] ldr r2, [r4, lr, lsl #2]	mov ecx,DWORD PTR [eax+12] mov eax,DWORD PTR [ecx+ebx*4]
cur_value -= coeff * value;	mufd f0, f2, f1 sufd f3, f3, f0	mulsd xmm1,xmm2 subsd xmm0,xmm1

<function prologue>	stmdb sp!, {r4, r5, lr}	push ebp
		mov ebp, esp
		and esp, 0xfffffff8
		sub esp, 0x10
		mov DWORD PTR [esp], esi
		mov DWORD PTR [esp+4], edi
		mov DWORD PTR [esp+8], ebx

<function epilogue>	ldmia sp!, {r4, r5, pc}	mov esi, DWORD PTR [esp] mov edi, DWORD PTR [esp+4] mov ebx, DWORD PTR [esp+8] mov esp, ebp pop ebp ret

SIMD

- MMX: Integer only. Renames the FP registers
- SSE: 8 new registers, each with 4 ints or 4 single-precision floats. Cache control.
- SSE2: Lots more data types
- SSE3: Horizontal operations
- SSE4: Unknown

Floating point

- 8087: Stack architecture
- P III, P4, Core: SSE, SSE2, SSE3, SSE4

Thread-level parallelism

- SMP
- Pentium Extreme Edition
- Core Duo
- Hyperthreading

SMP

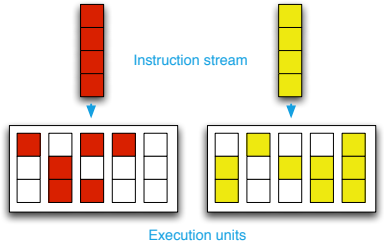

- P III, 4, Core: MESI (Athlon: MOESI)
- Bus arbitration via dedicated lines
- Interrupts: Whichever gets them

Hyperthreading

- Replicate register renaming, architectural registers
- Partition re-order buffer
- Share caches and execution units
- Relatively large performance gain

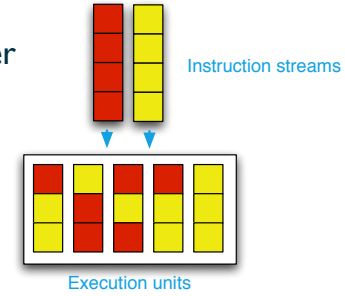

Performance

- "Speed demon" P4 was disappointing
- Focus on ILP / TLP
  - Dual-core designs
  - Hyperthreading

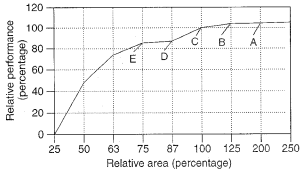

The future of IA-32

- EMT64
- Virtualisable
- Multicore

seuQnoit?s