



## Dynamic Scheduling in P6 (Pentium Pro, II, III)

- Q: How pipeline 1 to 17 byte 80x86 instructions?
- A: P6 doesn't pipeline 80x86 instructions
- P6 decode unit translates the Intel instructions into 72-bit microoperations (~ MIPS)
- Sends micro-operations to reorder buffer & reservation stations
- Many instructions translate to 1 to 4 micro-operations
- Complex 80x86 instructions are executed by a conventional microprogram (8K x 72 bits) that issues long sequences of micro-operations
- 14 clocks in total pipeline (~ 3 state machines)

Parameter	80x86	microops
Max. instructions issued/clock	3	6
Max. instr. complete exec./clock		5
Max. instr. committed/clock		3
Window (Instrs in reorder buffer)		40
Number of reservations stations		20
Number of rename registers		40
Number of integer functional units (FUs)		2
Number of floating point FUs		1
Number of SIMD floating point FUs		1
Number of memory FUs	1	load + 1 store

































#### Pentium 4

- Still translate from 80x86 to micro-ops
- P4 has better branch predictor, more functional units
- Instruction Cache holds micro-operations vs. 80x86 instructions
  - no decode stages of 80x86 on cache hit ("Trace Cache")
- Faster memory bus: 400 MHz v. 133 MHz
- Caches
  - Pentium III: L1-I 16KB, L1-D 16KB, L2 256 KB
  - Pentium 4: L1-I 12K uops, L1-D 8 KB, L2 256 KB
  - Block size: PIII 32B v. P4 128B; 128 v. 256 bits/clock
- Clock rates:
  - Pentium III 1 GHz v. Pentium IV 1.5 GHz
  - 14 stage pipeline vs. 24 stage pipeline

### **Trace Cache**

- IA-32 instructions are difficult to decode
- Conventional Instruction Cache
  - Provides instructions up to and including taken branch
- Trace cache, records uOps instead of x86 Ops
- Builds them into groups of six sequentially ordered uOps per line
  - Allows more ops per line
  - Avoids clock cycle to get to target of branch



- Multimedia instructions 128 bits wide vs. 64 bits wide => 144 new instructions
  - When used by programs??
  - Faster Floating Point: execute 2 64-bit Fl. Pt. Per clock
  - Memory FU: 1 128-bit load, 1 128-store /clock to MMX regs
- Using RAMBUS DRAM
  - Bandwidth faster, latency same as SDRAM
  - Cost 2X-3X vs. SDRAM
- ALUs operate at 2X clock rate for many ops
- Pipeline doesn't stall at this clock rate: uops replay
- Rename registers: 40 vs. 128; Window: 40 v. 126
- BTB: 512 vs. 4096 entries (Intel: 1/3 improvement)

refetch	Decode	Decode	Execute	Write-bac	:k					
	P5 /	Microarchit	tecture							
Fetch	Fetch	Decode	Decode	Decode	Rename	ROB Rd	Rdy/Sch	Dispatch	1 Execute	
				P6 Microa	architecture	2				
TC N	Ixt IP	тс	Fetch	Drive	Alloc	Rei	name	Queue	Schedule	
S	chedule	Schedule	Dispatch	Dispatch	Reg File	Reg File	Execute	Flags	Branch Ck	Drive
				netbul:	st microart	anceture				













# **IA-64 Registers**

- The integer registers designed to assist procedure calls using a register stack
  - Similar to SPARC's register windows.
  - Registers 0-31 are always accessible and addressed as 0-31
  - Registers 32-128 are used as a register stack and each procedure is allocated a set of registers (from 0 to 96)
  - The new register stack frame is created for a called procedure by renaming the registers in hardware;
  - a special register called the current frame pointer (CFM) points to the set of registers to be used by a given procedure
- 8 64-bit Branch registers used to hold branch destination addresses for indirect branches
- 64 1-bit predicate registers



# **IA-64 Registers**

- Both the integer and floating point registers support register rotation for registers 32-128.
- Register rotation eases the task of allocating registers in software pipelined loops
- Avoid the need for unrolling and for prologue and epilogue code for a software pipelined loop
  - Makes the SW-pipelining usable for loops with smaller numbers of iterations



Execution	Instruction	Instruction	Example
Unit Slot	type	Description	Instructions
l-unit	А	Integer ALU	add, subtract, and, or, cmp
	Ι	Non-ALU Int	shifts, bit tests, moves
M-unit	А	Integer ALU	add, subtract, and, or, cmp
	М	Mem access	Loads, stores for int/FP regs
F-unit	F	Floating point	Floating point instructions
B-unit	В	Branches	Conditional branches, calls
L+X	L+X	Extended	Extended immediates, stops

Templ	ate Examp	les			
	Template	Slot 0	Slot 1	Slot 2	]
	0	М	Ι	Ι	Stop bits
	1	М	Ι	Ι	
	2	М	Ι	Ι	]
	3	М	Ι	Ι	
					]
	28	М	F	В	
	29	М	F	В	
					-

# **Predication Support**

- Nearly all instructions are predicated
  - Conditional branches are predicated jumps!
- Compare/Test instructions set predicates
  - Ten different comparison tests + 2 predicate destinations
  - Written with result of comparison + complement

## **Speculation Support**

- All INT registers have a 1-bit NaT (Not A Thing)
  - This is a poison bit (as discussed earlier)
  - Speculative loads generate these
  - All other instructions propagate them
- Deferred exceptions
  - Nonspeculative exceptions receive a NAT as a source operand there is an unrecoverable exception
  - Chk.s instructions can detect and branch to recovery code

# **Memory Reference Support**

- Advanced Loads allow speculative memory references
  - Move loads ahead of potentially dependent stores
  - ALAT table is allocated with register destination + memory address
  - Stores associatively lookup the table when they execute
    - Invalidate ALAT entries with same memory address
- Before using the value of the advanced load
  - Explicit check is needed to see if ALAT entry is valid
  - If it fails, can re-load the value or perform cleanup operation

Itanium <sup>TM</sup> Mach	nine Characteristics
Frequency	800 MHZ
Transistor count	25.4M CPU; 295M L3
Process	0.18u CMOS, 6 metal layers
Package	Organic Land Grid Array
Machine width	6 instructions/clock (4 ALU/MM, 2 Ld/St, 2 FP, 3 Br)
Registers	14 ported 128 GR & 128 FR; 64 Predicates
Speculation	32 entry ALAT, Exception Deferral
Branch prediction	Multilevel 4-stage Prediction Hierarchy
FP compute bandwidth	3.2 GFlops (DP/EP); 6.4 GFlops (SP)
Memory $\rightarrow$ FP bandwidth	4 DP (8 SP) operands/clock
Virtual memory support	64 entry ITLB, 32/96 2-level DTLB, VHPT
L2/L1 cache	Dual ported 96K unified & 16KD; 16KI
L2/L1 latency	6 / 2 clocks
L3 cache	4MB, 4-way s.a., BW of 12.8 GB/sec;
System bus	2.1 GB/sec; 4-way glueless MP; scalable to large (512+ proc) systems





## Itanium processor 10-stage pipeline

- Front-end (stages IPG, Fetch, and Rotate): prefetches up to 32 bytes per clock (2 bundles) into a prefetch buffer, which can hold up to 8 bundles (24 instructions)
  - Branch prediction is done using a multilevel adaptive predictor like P6 microarchitecture
- Instruction delivery (stages EXP and REN): distributes up to 6 instructions to the 9 functional units
  - Implements registers renaming for both rotation and register stacking.

### Itanium processor 10-stage pipeline

- Operand delivery (WLD and REG):
  - Accesses register file
  - Performs register bypassing
  - Accesses and updates a register scoreboard
    - Scoreboard used to detect when individual instructions can proceed, so that a stall of 1 instruction in a bundle need not cause the entire bundle to stall
  - Checks predicate dependences.

## Itanium processor 10-stage pipeline

- Execution (EXE, DET, and WRB)
  - Executes instructions through ALUs and load/store units
  - Detects exceptions and posts NaTs
  - Retires instructions and performs write-back
  - Deferred exception handling via poison bits (NaTs)
- Predicate Delivery
  - Predicates generated in EXE delivered in DET and feed into retirement, branch execution, dependency detect
  - All instructions read operands and execute
  - Canceled at retirement



# **IBM PowerPC 620**

- Core design reused for POWER3, POWER4, POWER5: IBM servers
- Core used in game machines: PS2, XBOX
- 4-way superscalar
- OOO BTB (BTAC), BHT.
- Distributed rsvn stations
- Register renaming with separate rename buffers

PowerPC 620 pipel	ne
Fetch stage Instruction buffer (8) Dispatch stage	
Reservation stations (6) Execute stage(s) Completion buffer (16) Complete stage Writeback stage	XSU0 XSU1 MC-FXU LSU FPU FPU FPU FPU FPU FPU FPU FPU FPU FP



