Computer Architecture

Spring 2016

Lecture 15: Dynamic Scheduling - Case Study

Shuai Wang

Department of Computer Science and Technology

Nanjing University

[Slides adapted from CS 246, Harvard University]

Dynamic Scheduling in P6 (Pentium Pro, II, III, M, Core, Core2Duo, Core i7)

Q: How pipeline 1 to 17 byte 80x86 instructions?

- P6 doesn’t pipeline 80x86 instructions
- P6 decode unit translates the Intel instructions into 72-bit micro-operations (~ 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)

What is a uop?

- **Small two-operand instruction - Very RISC like.**

  IA-32 instruction

  add (eax),(ebx) MEM(eax) <- MEM(eax) + MEM(ebx)

  Uop decomposition:

ld guop0,(eax)	guop0 <- MEM(eax)
ld guop1,(ebx)	guop1 <- MEM(ebx)
add guop0,guop1	guop0 <- guop0 + guop1
sta eax	
std guop0	MEM(eax) <- guop0

P6 Pipeline

- 8 stages are used for in-order instruction fetch, decode, and issue
  - Takes 1 clock cycle to determine length of 80x86 instructions + 2 more to create the micro-operations (uops)
- 3 stages are used for out-of-order execution in one of 5 separate functional units
- 3 stages are used for instruction commit

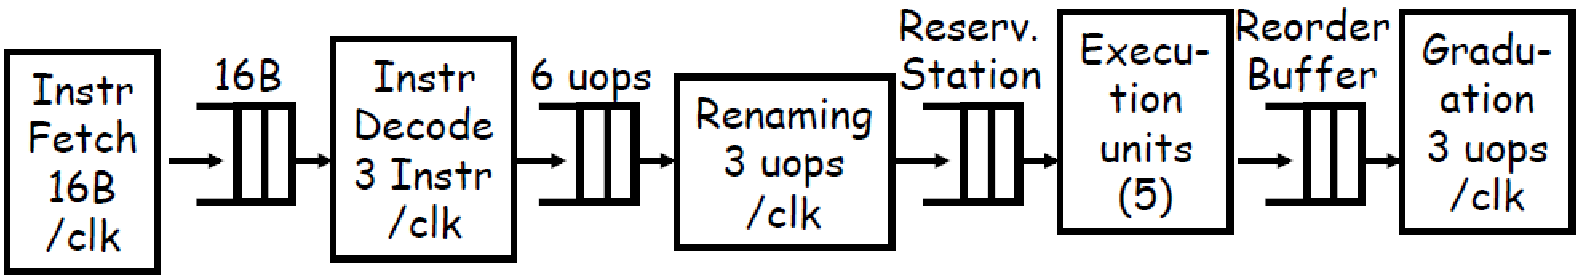

Pentium III Overview

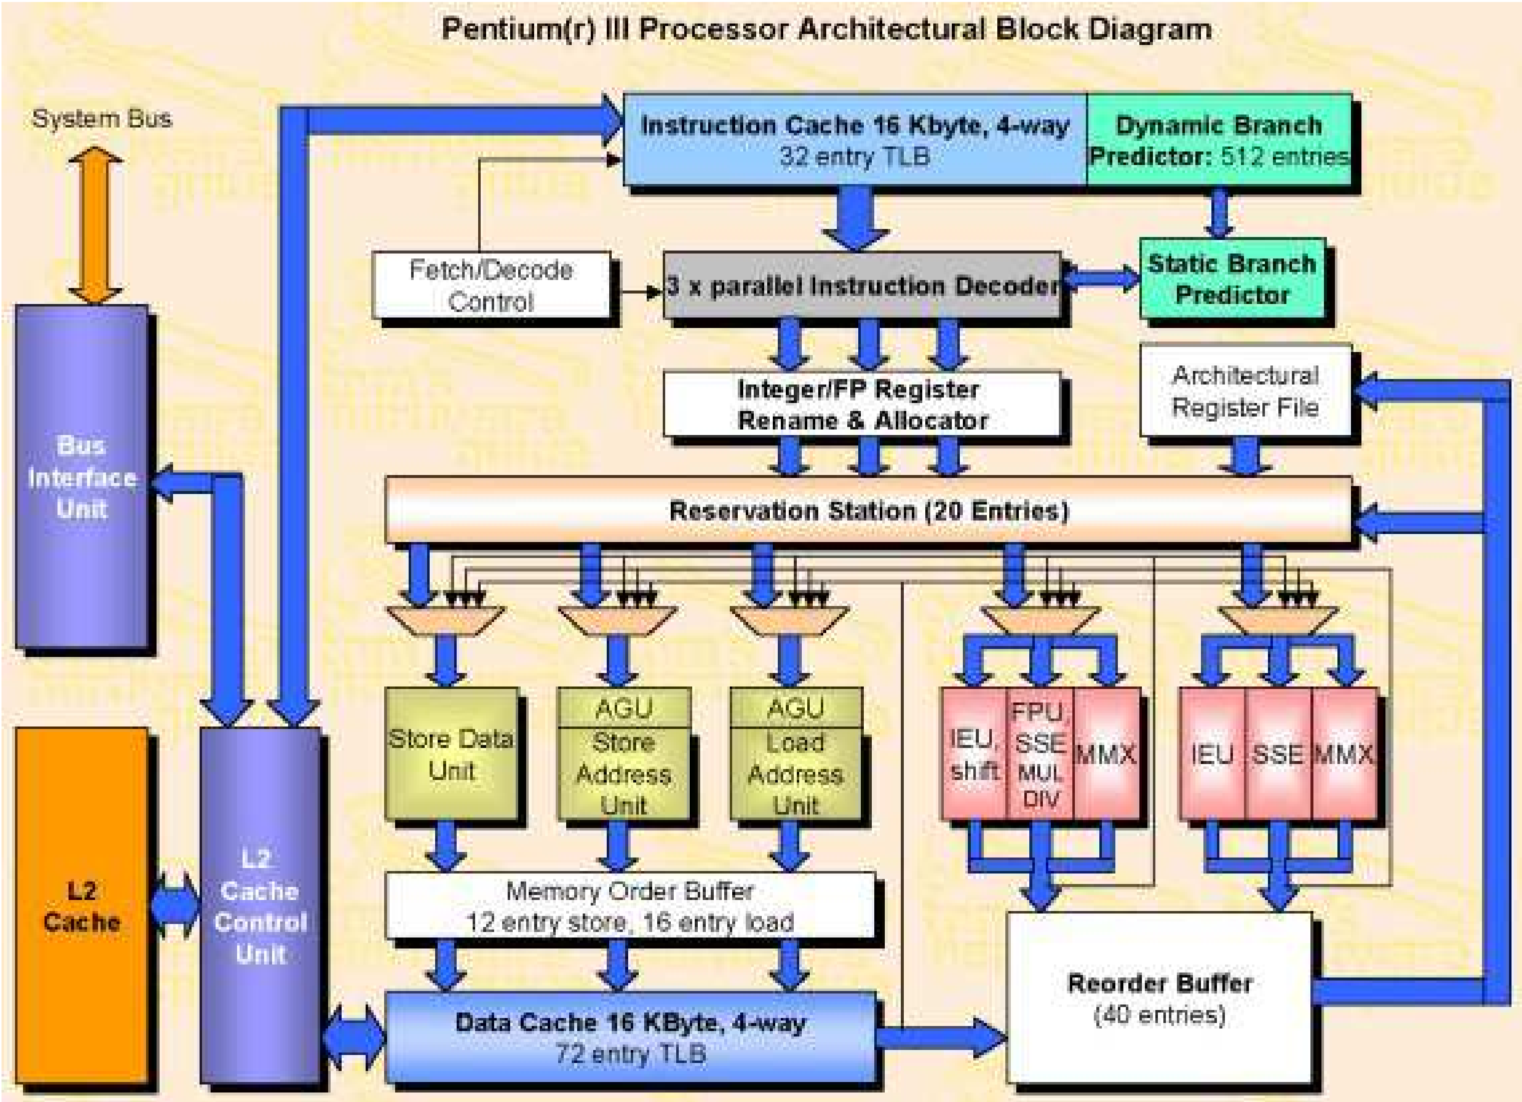

Pentium Pro Design Methodology

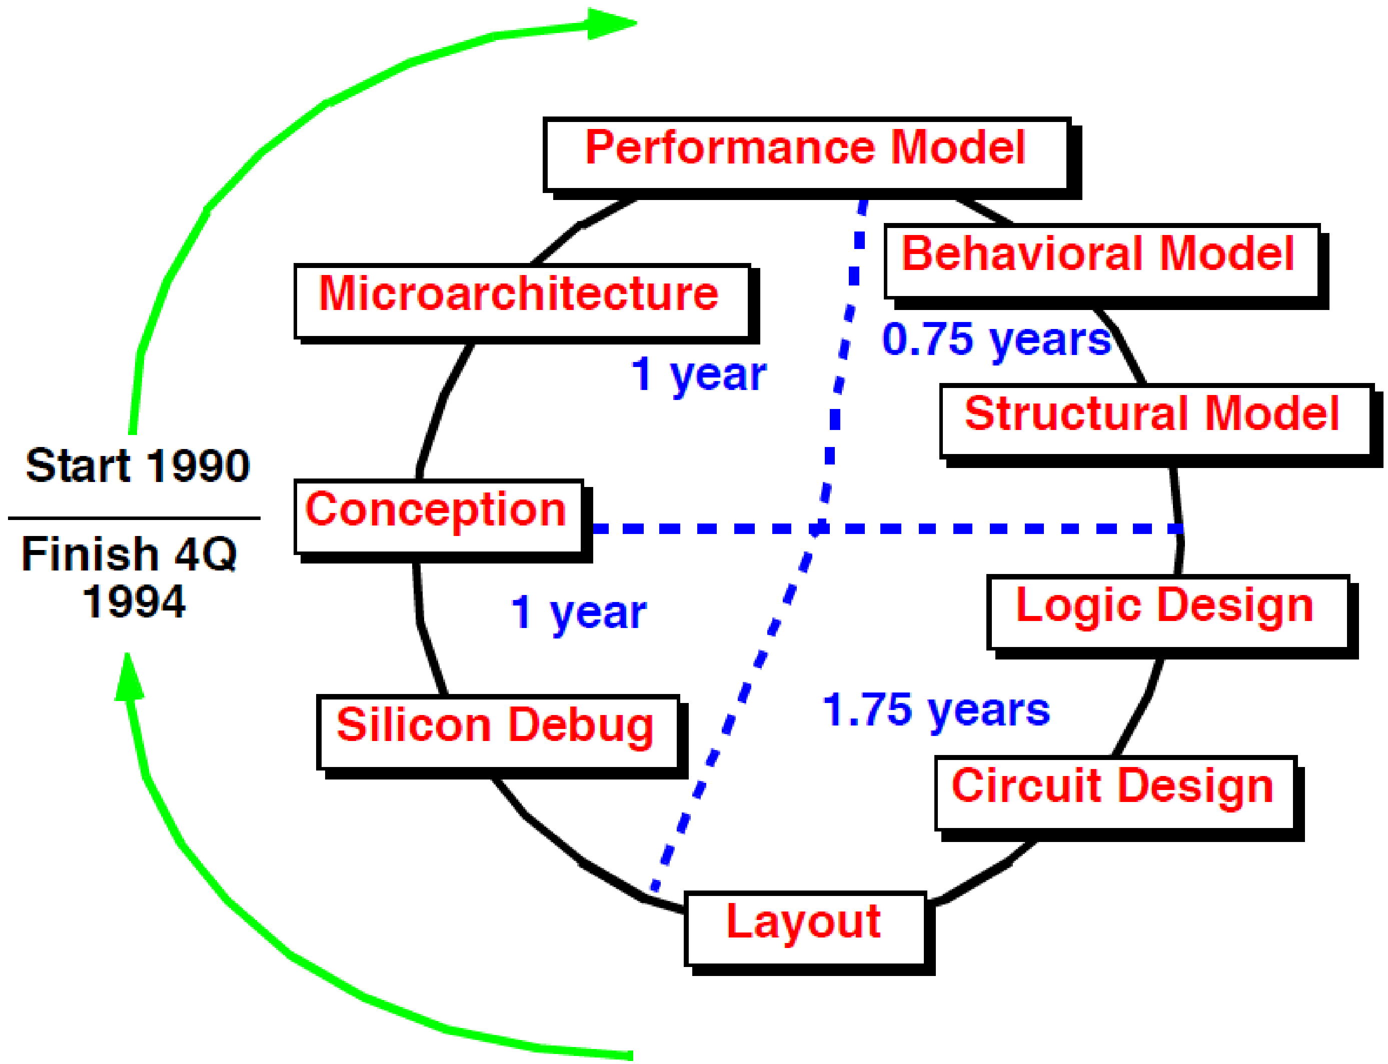

Retrospective

- Most commercially successful microarchitecture in history
- Evolution
  - Pentium II/III, Xeon, etc.
    - Derivatives with on-chip L2, ISA extensions, etc.
  - Replaced by Pentium 4 as flagship in 2001
    - High frequency, deep pipeline, extreme speculation
  - Resurfaced as Pentium M in 2003
    - Initially a response to Transmeta in laptop market
    - Pentium 4 derivative (90nm Prescott) delayed, slow, hot
  - Core Duo, Core 2 Duo, Core i7 replaced Pentium 4

Microarchitectural Updates

- Pentium M (Banias), Core Duo (Yonah)
  - Micro-op fusion (also in AMD K7/K8)
    - Multiple uops in one: (add eax,[mem] => ld/alu), sta/std
    - These uops decode/dispatch/commit once, issue twice
  - Better branch prediction
    - Loop count predictor
    - Indirect branch predictor
  - Slightly deeper pipeline (12 stages)
    - Extra decode stage for micro-op fusion
    - Extra stage between issue and execute (for RS/PLRAM read)
  - Data-capture reservation station (payload RAM)
    - Clock gated for 32 (int) , 64 (fp), and 128 (SSE) operands

Microarchitectural Updates

- Core 2 Duo (Merom)
  - 64-bit ISA from AMD K8
  - Macro-op fusion
    - Merge uops from two x86 ops
    - E.g. cmp, jne => cmpjne
  - 4-wide decoder
    - Peak x86 decode throughput is 5 due to macro-op fusion
  - Loop buffer
  - Loops up to 18 uops avoid fetch/decode overhead
  - Even deeper pipeline (14 stages)
  - Larger reservation station (32), instruction window (96)

Microarchitectural Updates

- Nehalem (Core i7/i5/i3)
  - RS size 36, ROB 128
  - Loop cache up to 28 uops
  - L2 branch predictor
  - L2 TLB
  - I$ and D$ now 32K, L2 back to 256K, inclusive L3 up to 8M
  - Simultaneous multithreading
  - RAS now renamed (repaired)
  - 6 issue, 48 load buffers, 32 store buffers
  - New system interface (QPI) – finally dropped front-side bus
  - Integrated memory controller (up to 3 channels)
  - New STTNI instructions for string/text handling

Pentium 4

- Still translate from 80x86 to micro-ops
- P4 has better branch predictor, more FUs
- 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: L1I 16KB, L1D 16KB, L2 256 KB
  - Pentium 4: L1I 12K uops, L1D 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

Pentium III vs. Pentium4

	Pentium III	Pentium 4
Peak Clock	450 MHz (Xeon)	2.4 (4.8 internal) GHz
Pipeline stages	14	22
Branch Prediction	512 entry local	2K entry hybrid
BTB	512 entries	4K entries
I$	16KB	64 KB t$ + 8KB I$
D$	16KB	8kb
L2	512KB-2MB	256KB-2MB
Fetch Width	16 bytes	3 μops (16 bytes on miss)
Rename/Retire Width	3 μops	3 μops
Execute Width	5 μops	7 μops (X2)
Reservation Stations	20	60
ROB Size	40	128
Register Renaming	P6-style	MIPS R10K-style
Memory disambiguation	Conservative	Predictor-Based
Anything else?	No	Multithreading

Reading Assignment

- K. C. Yeager. "***The MIPS R10000 Superscalar Microprocessor***." IEEE Micro, pp. 28-40, 16(2), April 1996

- B. Sinharoy et. Al. "***Power5 System Microarchitecture***." IBM Journal of Research and Development, pp. 505 – 521, 49(4-5), July 2005