

1. PowerPC 970FX (G5) Overview

Frontend/Backend

- RISC-like frontend, but inst set old and baroque
 - Many instructions microcoded (cracked into multiple instruction)
- VLIW-like backend
 - 5 inst issued in 'group', one of which must be branch
 - Group tracked together, but not explicitly superscalar as in true VLIW
- 64 bit hardware
 - MacOS 10 32 bit
 - Linux 64 bit

Architecture Summary

- 32KB L1 data cache
- 64KB 2-way L1 inst cache
- 512KB 8-way L2 cache
- 1024 entry, 4-way TLB
- Two MAC FPU units (4 FLOP/cycle, 6 cycle lat)
- 1 Vector permute unit
- 1 VALU (8 sFLOP/cycle)
- 8 hardware prefetch streams
- 4 software vector prefetch streams (use 4 of hardware streams)
- 32 logical ireg/freg/vreg, 80 physical

2. PPC970FX Architecture Summary

- Out-of-order exec up 10 ops in 10 pipes:
 - Two load/store ops
 - Two fixed-pt reg-reg ops
 - Two fp ops
 - One branch
 - One condition reg op
 - One vector permute op
 - One vector ALU op
- Register renaming on almost all regs
- Speculative superscalar execution
- Up to 215 inst in flight at once
 - 16 inst in inst fetch unit
 - 32 inst buffered in inst decode unit
 - 35 inst in 3 decode pipe stges and four dispatch buffers
 - 100 inst in inner-core (after dispatch)
 - 32 stores in store queue (STQ)
- Out-of-order and speculative execution of loads

3. PPC970FX Caches

L1 Data Cache

- 32KB, 2-way assoc
- Write-through
- LRU replacement
- 128-byte linesize
 - Loaded in 4 32-byte "beats"
 - loaded beat fetched first
- clocked 1:1 with CPU:
 - 2-cycle ld-use penalty FXU ld
 - 4-cycle ld-use penalty FPU ld
 - 3-cycle ld-use penalty VPERM ld
 - 4-cycle ld-use penalty VALU ld

L2 Combined Cache

- 512KB, 8-way assoc
- LRU replacement
- 128-byte linesize
- inclusive
- 8 64-byte store queue between L1 & L2
- Up to 8 hardware directed prefetch streams into L2 and L1
 - recognizes increasing/decreasing sequential access

4. PPC970FX Prefetching – Data block touch

- Vector versions deprecated
- `dcbt ra, rb, imm`
- Prefetch for read
- Effective Address (EA) $ra/0 + rb$
- Stream variants pref even if cache-resident
- 2 & 4-bit imm forms

2-bit immediate:

- 0b00: pref line containing EA
- 0b01: pref stream inc from EA
- 0b10: reserved
- 0b11: pref stream dec from EA

4-bit immediate

- 0b0000: pref line containing EA
- 0b0001: pref stream inc from EA
- 0b0011: pref stream dec frm EA
- 0b1000: pref according to hint from bits 57-63 of EA:
 - Bits 0-56: High order 57 bits of EA to fetch (low=0)
 - Bit 57: Direction (0-forward, 1-backward)
 - Bit 58: 0-No info; 1-unlimited, reused stream
 - Bit 59: reserved
 - Bits 60-63: Hardware stream ID to use

5. PPC970FX Vector Units

- Need to turn on 'Java' mode to make IEEE compliant
 - Denorms handled in software when java mode on, not handled at all otherwise
 - Unsure effect on performance, increased latency in G4
- Shares load/store units with FPU
- 8-cycle latency on VFPU
- 2-cycle latency on VPERM
- Extra cycle to transfer between VPERM and VALU