ECE 571 – Advanced Microprocessor-Based Design Lecture 8

Vince Weaver

http://www.eece.maine.edu/~vweaver

vincent.weaver@maine.edu

25 September 2014

Announcements

- HW2 was posted, due Tuesday

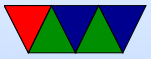

Misc

- Cortex-M7 released, big change from Cortex-M4 is it went Superscalar with 2-issue. (M4 1-issue 3-stage pipeline. M7= 2-issue 6-stage pipeline)
- Ask about flush pipeline on context switch. Probably, what triggers it is virtual memory change and TLB flush

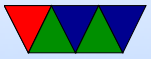

Long Pipelines

Since these are Out-of-Order processors these pipelines are more of a summary than an actual depiction of what hardware does.

Pentium III 10–stage Pipeline

1	2	3	4	5	6	7	8	9	10
Fetch		Decode	Decode		Rename	ROB	Sched	Dispatch	Exec

Pentium 4 20–stage Pipeline

1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20
NextIP	TraceCache		Fetch	Drive	Alloc	Rename		Que	Schedule			Dispatch		Reg File		Exe	Flags	Br Ck	Drive

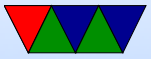

HW#1 Review

- 1st question, instructions and cycles.
  Instructions should be deterministic, almost is
  Cycles is not deterministic.
  Does re-ordering link order matter? Yes, but possibly lost in noise.

- 2nd question. Should have been called "sampling and profiling", not just sampling.
  Did all methods give same results? Mostly.
  Timing: gprof and perf about the same, valgrind much

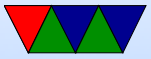

slower.

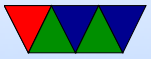

HW#1 Review

cycles:				
0.55	6b4:		sub	$0x4,%rax
0.18			sub	$0x1,%edx
0.18			cmp	%r8d,%ecx
		↑	jl	62a
	6c4:		mov	(%rax),%ecx
0.55			lea	(%rcx,%r13,1),%esi
0.74			movzbl	(%r15,%rsi,1),%esi
4.04			cmp	%r9d,%esi
		↑	je	6a0
0.37		↓	js	6da
0.55			mov	%edx,%ecx

movzbl = move byte zero extended to long

address 0x402d1a

use `addr2line` and compile with -g to find code in question:

n = ((Int32)block[ptr[unHi]+d]) - med;

If force perf to use low-skid PEBS events we can see skid is a problem, and the actual results are as such:

cycles:pp

	6b4:		sub	$0x4,%rax
0.72			sub	$0x1,%edx
0.18			cmp	%r8d,%ecx
		↑	jl	62a
0.90	6c4:		mov	(%rax),%ecx
0.72			lea	(%rcx,%r13,1),%esi
2.69			movzbl	(%r15,%rsi,1),%esi
0.54			cmp	%r9d,%esi
		↑	je	6a0
1.44		↓	js	6da
0.72			mov	%edx,%ecx

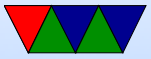

CPU Power and Energy

- Became a trendy thing to research in 1999-2002 timeframe.
- Before that usually concern was with performance.
- These days energy results are often reported as a core part of any architectural proposal, not as a separate issue.
- The results discussed here are academic and may or may not be implemented in actual chips.

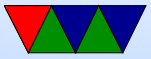

AMD Bulldozer Die Shot

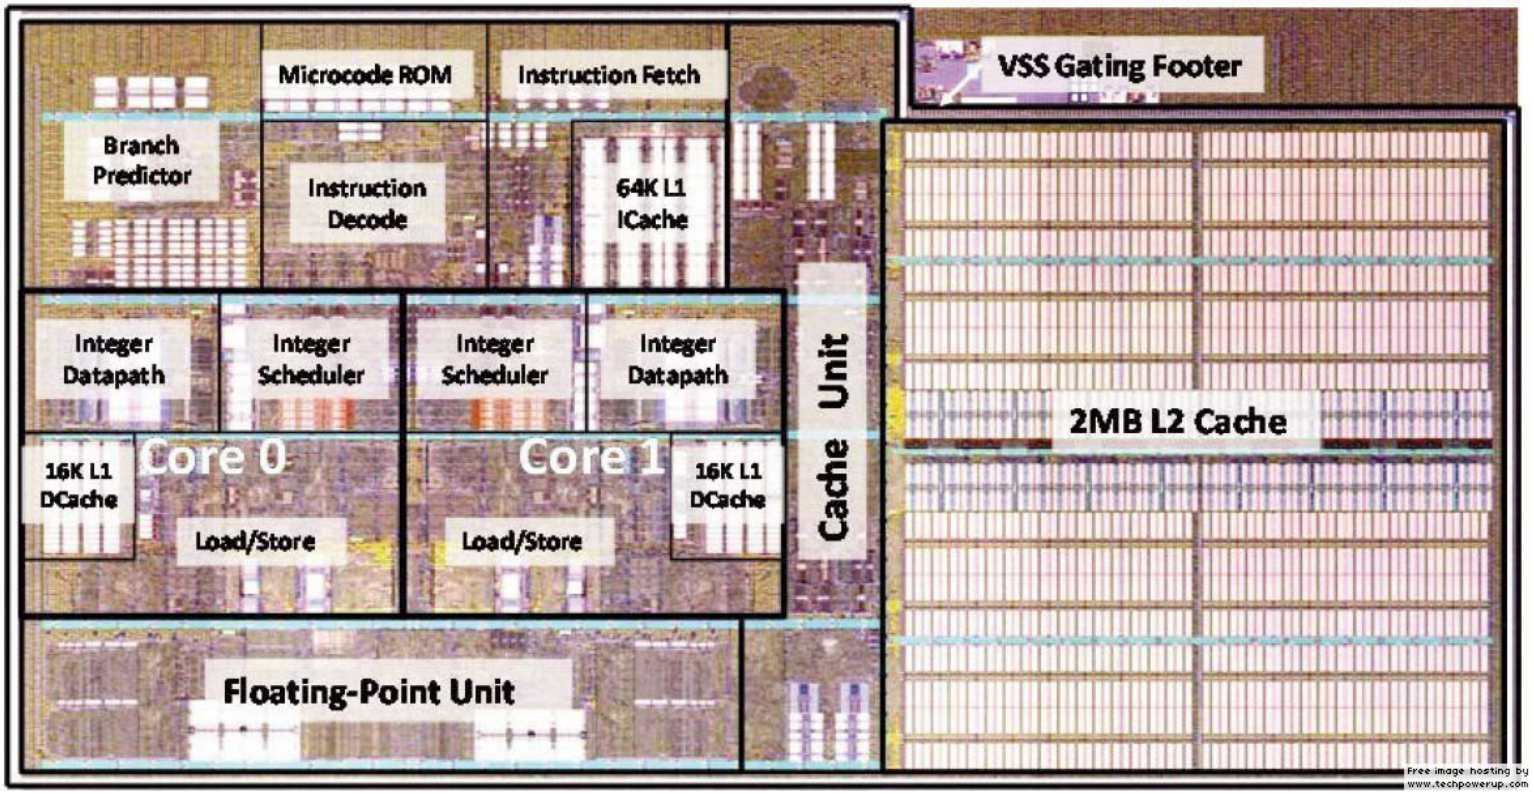

Note which structures are big, using static power.

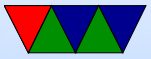

CPU Power Breakdown

From Fan, Tang, Huan, Gao (ISLPED'05), Chinese Godson MIPS CPU

They gave numbers, but unclear of workload, if static or dynamic, etc.

- Cache 36%
- TLB 13%
- FALU 10%
- ROQueue 7%
- FMUL 6%

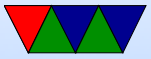

- Float reg 5%
- Gen reg 5%
- MUL 2%
- MCUControl 2%
- ALU 1%
- Other 13%

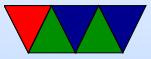

Thermal Concerns Too

Power density exceed hot plate, approaching rocket nozzle

TODO: Find the Intel cite for this statement.

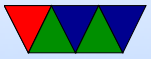

Methodologies Used in These Papers

It varies, but many of these are from simulations (sometimes validated). Anything from SPICE to "cycle-accurate" simulators.

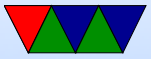

Clock Generation

- Driving high-frequency load against capacitance, trying to keep whole chip in sync.
- Extreme Case: Alpha 21264 H-tree, 32% of power?
- Half-frequency clocks (on both edge, so clock run half as fast) (Mudge 2001)
- Asynchronous
- Locally Asynchronous (Divide to multiple clock domains)

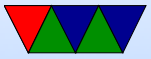

DVFS and other CPU Power/Energy Saving Methods

- A lot of related work
- Will focus on actual implementations rather than academic papers this time

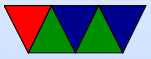

DVFS

- Voltage planes – on CMP might share voltage planes so have to scale multiple processors at a time
- DC to DC converter, programmable.
- Phase-Locked Loops. Orders of ms to change. Multiplier of some crystal frequency.
- Senger et al ISCAS 2006 lists some alternatives. Two phase locked loops? High frequency loop and have programmable divider?

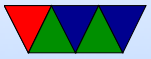

- Often takes time, on order of milliseconds, to switch frequency. Switching voltage can be done with less hassle.

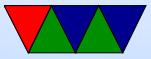

Adaptive Body Biasing

- Related to but not always considered part of DVFS
- Control voltage applied to body
- Change the threshold voltage
- Reduces leakage but slows performance