CS 152 Computer Architecture and Engineering CS252 Graduate Computer Architecture

Lecture 13 –VLIW

Chris Fletcher Electrical Engineering and Computer Sciences University of California at Berkeley

https://cwfletcher.github.io/
http://inst.eecs.berkeley.edu/~cs152

Last Time in Lecture 12

- Branch prediction
 - temporal, history of a single branch
 - spatial, based on path through multiple branches
- Branch History Table (BHT) vs. Branch History Buffer (BTB)
 - tradeoff in capacity versus latency
- Return-Address Stack (RAS)
 - specialized structure to predict subroutine return addresses
- Advanced branch prediction structures
 - Perceptron
 - TAGE

Superscalar Control Logic Scaling



- Each issued instruction must somehow check against W*L instructions, i.e., growth in hardware \$\alpha\$ W*(W*L)
- For in-order machines, L is related to pipeline latencies and check is done during issue (interlocks or scoreboard)
- For out-of-order machines, L also includes time spent in instruction buffers (instruction window or ROB), and check is done by broadcasting tags to waiting instructions at write back (completion)
- As W increases, larger instruction window is needed to find enough parallelism to keep machine busy => greater L

=> Out-of-order control logic grows faster than W^2 (~ W^3)

Out-of-Order Control Complexity:



Sequential ISA Bottleneck



VLIW: Very Long Instruction Word



- Multiple operations packed into one instruction
- Each operation slot is for a fixed function
- Constant operation latencies are specified
- Architecture requires guarantee of:
 - Parallelism within an instruction => no cross-operation RAW check
 - No data use before data ready => no data interlocks

Early VLIW Machines

FPS AP120B (1976)

- scientific attached array processor
- first commercial wide instruction machine
- hand-coded vector math libraries using software pipelining and loop unrolling
- Multiflow Trace (1987)
 - commercialization of ideas from Fisher's Yale group including "trace scheduling"
 - available in configurations with 7, 14, or 28 operations/instruction
 - 28 operations packed into a 1024-bit instruction word
- Cydrome Cydra-5 (1987)
 - 7 operations encoded in 256-bit instruction word
 - rotating register file

VLIW Compiler Responsibilities

- Schedule operations to maximize parallel execution
- Guarantees intra-instruction parallelism

- Schedule to avoid data hazards (no interlocks)
 - Typically separates operations with explicit NOPs

Loop Execution



How many FP ops/cycle?

1 fadd / 8 cycles = 0.125

Loop Unrolling

for (i=0; i<N; i++)

B[i] = A[i] + C;

Unroll inner loop to perform 4 iterations at once

for (i=0; i<N; i+=4)

{

}

B[i] = A[i] + C;

B[i+1] = A[i+1] + C;

B[i+2] = A[i+2] + C;

```
B[i+3] = A[i+3] + C;
```

Need to handle values of N that are not multiples of unrolling factor with final cleanup loop

Scheduling Loop Unrolled Code



How many FLOPS/cycle?

4 fadds / 11 cycles = 0.36

Software Pipelining

Unroll 4 ways first			Int1	Int 2	M1	M2	FP+	FPx
loop: fld f1, 0(x1)		C			fld f1			
fld f2, 8(x1)					fld f2			
fld f3, 16(x1)					fld f3			
fld f4, 24(x1)	prolog <	prolog	add x	1	fld f4			
add x1, 32					fld f1		fadd f5	
fadd f5, f0, f1					fld f2		fadd f6	
fadd f6, f0, f2					fld f3		fadd f7	
fadd f7, f0, f3				add x	1	fld f4		fadd f8
fadd f8, f0, f4		rate loop:			fld f1	fsd f5	fadd f5	
fsd f5, 0(x2)					fld f2	fsd f6	fadd f6	
fsd f6, 8(x2)				add x2	fld f3	fsd f7	fadd f7	
fsd f7, 16(x2)			add x	1bne	fld f4	fsd f8	fadd f8	
add x2, 32						fsd f5	fadd f5	
fsd f8, -8(x2)	epilog					fsd f6	fadd f6	
bne x1, x3, loop		^{og} \prec		add x2		fsd f7	fadd f7	
				bne		fsd f8	fadd f8	
How many FLOPS/cycle?						fsd f5		
4 fadds / 4 cyc	les =	Τ ~			·		·	

Software pipelining

How do we write the loop?

Loop: fld **f0**,0(x1) fadd **f4**,**f0**,f2 fsd **f4**,0(x1)

addi x1,x1,-8

bne x1,x2,Loop



SW pipelined, main loop:

- Loop': fsd f4,16(x1) fadd f4,f0,f2 fld f0,0(x1) addi x1,x1,-8
 - bne x1,x2,Loop'

Startup/Wind-down code

	Startup:				
Basic idea:		fld	f0, 0(x1)	//	ld, iter 0
Add some code before/after main loop to ensure		fadd	f4,f0, f2	//	add, iter 0
that		fld	f0 ,-8(x1)	//	ld, iter 1
		addi	x1,x1,-1 6		
Each logical iteration has a ld, add, sub	Loop':				
		fsd	f4, 16(x1)	//	st, iter 0
		fadd	f4,f0, f2	//	add, iter 1
		fld	f0, 0(x1)	//	ld, iter 2
		addi	x1,x1,- 8		
		bne	<pre>x1,x2,Loop'</pre>		
	Winddown:				
		fsd	f4 ,16(x1)	11	st, iter N-1
		fadd	f4,f0, f2	11	add, iter N
		fsd	f4, 0(x1)	//	st, iter N

Software Pipelining vs. Loop Unrolling



Problems with "Classic" VLIW

- Object-code compatibility
 - have to recompile all code for every machine, even for two machines in same generation
- Object code size
 - instruction padding wastes instruction memory/cache
 - loop unrolling/software pipelining replicates code
- Scheduling variable latency memory operations
 - caches and/or memory bank conflicts impose statically unpredictable variability
- Knowing branch probabilities
 - Profiling requires an significant extra step in build process
- Scheduling for statically unpredictable branches
 - optimal schedule varies with branch path

VLIW Instruction Encoding



- Schemes to reduce effect of unused fields
 - Compressed format in memory, expand on I-cache refill
 - used in Multiflow Trace
 - introduces instruction addressing challenge
 - Mark parallel groups
 - used in TMS320C6x DSPs, Intel IA-64
 - Provide a single-op VLIW instruction
 - Cydra-5 UniOp instructions

Intel Itanium, EPIC IA-64

- EPIC is the style of architecture (cf. CISC, RISC)
 - Explicitly Parallel Instruction Computing (really just VLIW)
- IA-64 is Intel's chosen ISA (cf. x86, MIPS)
 - IA-64 = Intel Architecture 64-bit
 - An object-code-compatible VLIW
- Merced was first Itanium implementation (cf. 8086)
 - First customer shipment expected 1997 (actually 2001)
 - McKinley, second implementation shipped in 2002
 - Recent version, Poulson, eight cores, 32nm, announced 2011

Eight Core Itanium "Poulson" [Intel 2011]



- 8 cores
- 1-cycle 16KB L1 I&D caches
- 9-cycle 512KB L2 I-cache
- 8-cycle 256KB L2 D-cache
- 32 MB shared L3 cache
- 544mm² in 32nm CMOS
- Over 3 billion transistors

- Cores are 2-way multithreaded
- 6 instruction/cycle fetch
 - Two 128-bit bundles
- Up to 12 insts/cycle execute

IA-64 Instruction Format



- Template bits describe grouping of these instructions with others in adjacent bundles
- Each group contains instructions that can execute in parallel



IA-64 Registers

- 128 General Purpose 64-bit Integer Registers
- 128 General Purpose 64/80-bit Floating Point Registers
- 64 1-bit Predicate Registers
- GPRs "rotate" to reduce code size for software pipelined loops
 - Rotation is a simple form of register renaming allowing one instruction to address different physical registers on each iteration

Intel Kills Itanium

- Donald Knuth " ... Itanium approach that was supposed to be so terrific—until it turned out that the wished-for compilers were basically impossible to write."
- "Intel officially announced the end of life and product discontinuance of the Itanium CPU family on January 30th, 2019", Wikipedia

What if there are no loops?



- Branches limit basic block size in control-flow intensive irregular code
- Difficult to find ILP in individual basic blocks

Trace Scheduling [Fisher,Ellis]



- Pick string of basic blocks, a *trace*, that represents most frequent branch path
- Use <u>profiling feedback</u> or compiler heuristics to find common branch paths
- Schedule whole "trace" at once
- Add fixup code to cope with branches jumping out of trace

Rotating Register Files

Problems: Scheduled loops require lots of registers, Lots of duplicated code in prolog, epilog

Solution: Allocate new set of registers for each loop iteration



Rotating Register File



Rotating Register Base (RRB) register points to base of current register set. Value added on to logical register specifier to give physical register number. Usually, split into rotating and nonrotating registers.

Rotating Register File (Previous Loop Example)

Three cycle load latency encoded as difference of 3 in register specifier number (f4 - f1 = 3) Four cycle fadd latency encoded as difference of 4 in register specifier number (f9 – f5 = 4)

/		¥	
ld f1, ()	fadd f5, f4,	sd f9, ()	bloop

ld P9, ()	fadd P13, P12,	sd P17, ()	bloop	RRB=8
ld P8, ()	fedd P12, P11,	sd P16, ()	bloop	RRB=7
ld P7, ()	fadd P11, P10,	sd P15, ()	bloop	RRB=6
ld P6, ()	fadd P10, P9,	sd P14, ()	bloop	RRB=5
ld P5, ()	fadd P9, P8,	sd P13, ()	bloop	RRB=4
ld P4, ()	fadd P8, P7,	sd P12, ()	bloop	RRB=3
ld P3, ()	fadd P7, P6,	sd P11, ()	bloop	RRB=2
ld P2, ()	fadd P6, P5,	sd P10, ()	bloop	RRB=1
		· · · · · · · · · · · · · · · · · · ·		

CS252



IA-64 Predicated Execution

Problem: Mispredicted branches limit ILP

Solution: Eliminate hard to predict branches with predicated execution

- Almost all IA-64 instructions can be executed conditionally under predicate
- Instruction becomes NOP if predicate register false



Four basic blocks

Warning: Complicates bypassing!

IA-64 Speculative Execution

Problem: Branches restrict compiler code motion

Solution: Speculative operations that don't cause exceptions



Particularly useful for scheduling long latency loads early

CS252

IA-64 Data Speculation

Problem: Possible memory hazards limit code scheduling **Solution**: Hardware to check pointer hazards



Requires associative hardware in address check table

CS252

Limits of Static Scheduling

- Statically unpredictable branches
- Variable memory latency (unpredictable cache misses)
- Code size explosion
- Compiler complexity
- Despite several attempts, VLIW has failed in general-purpose computing arena (so far).
 - More complex VLIW architectures are close to in-order superscalar in complexity, no real advantage on large complex apps.
- Successful in embedded DSP market
 - Simpler VLIWs with more constrained environment, friendlier code.

Acknowledgements

- This course is partly inspired by previous MIT 6.823 and Berkeley CS252 computer architecture courses created by my collaborators and colleagues:
 - Arvind (MIT)
 - Joel Emer (Intel/MIT)
 - James Hoe (CMU)
 - John Kubiatowicz (UCB)
 - David Patterson (UCB)
 - Krste Asanovic (UCB)
 - Sophia Shao (UCB)