VLIW/EPIC: Statically Scheduled ILP

Krste Asanovic
Laboratory for Computer Science
Massachusetts Institute of Technology

http://www.csg.lcs.mit.edu/6.823

Little's Law

Parallelism = Throughput * Latency

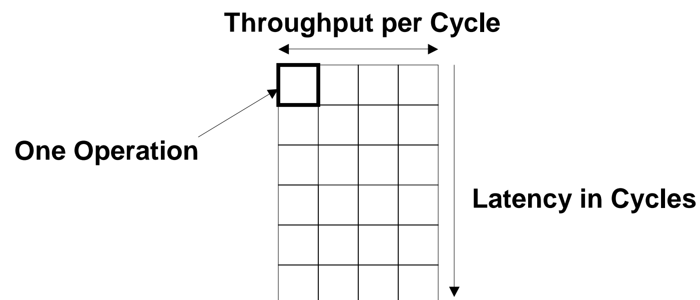

Example Pipelined ILP Machine

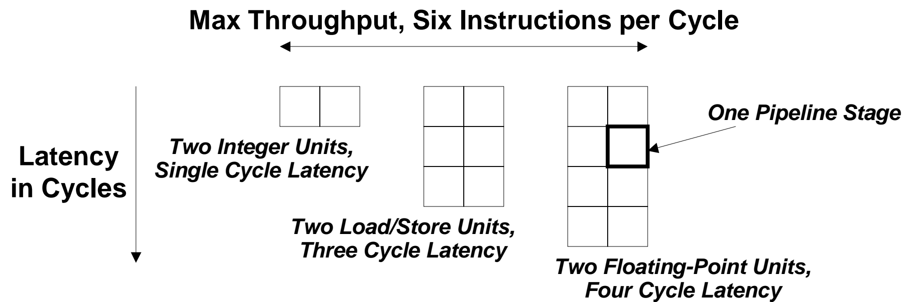

- How much parallelism (i.e., how many independent instructions) required to keep machine pipelines busy?

Superscalar Control Logic Scaling

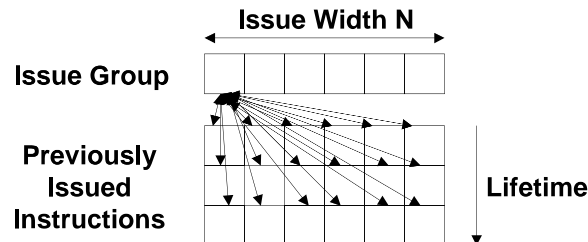

- Number of interlock checks and/or tag compares grows as N*(N*L) where L is lifetime of instructions in machine
  - Each of N instructions issued must check against N*L instructions in flight
- For in-order machines, lifetime L is related to pipeline latencies
- For out-of-order machines, L also includes time spent in instruction buffers (instruction window)
- As N increases, need larger instruction window to find enough parallelism to keep machine busy => greater lifetime L

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

ILP Control Complexity: MIPS R10000

Control Logic

Sequential ISA Bottleneck

Sequential source code

```
a = foo(b);
for (i=0, i<
```

Superscalar compiler

Find independent operations

Schedule operations

Sequential machine code

Superscalar processor

Check instruction dependencies

Schedule execution

VLIW: Very Long Instruction Word

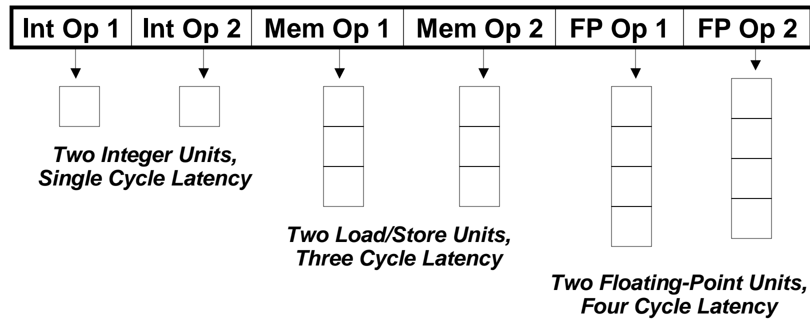

- Compiler schedules parallel execution
- Multiple parallel operations packed into one long instruction word
- Compiler must avoid data hazards (no interlocks)

Early VLIW Machines

- FPS 120AB (scientific attached array processor)
  - first commercial wide instruction machine
  - hand-coded vector math libraries using software pipelining and loop unrolling
- Multiflow Trace
  - 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
  - 7 operations encoded in 256-bit instruction word
  - rotating register file

Loop Execution

```
for (i=0; i<N; i++)
    B[i] = A[i] + C;
```

Compile

```
loop:  ld f1, 0(r1)
       add r1, 8
       fadd f2, f0, f1
       sd f2, 0(r2)
       add r2, 8
       bne r1, r3, loop
```

Schedule

loop:

Int1	Int 2	M1	M2	FP+	FPx

Loop Execution

```
for (i=0; i<N; i++)
    B[i] = A[i] + C;
```

Compile

```
loop:  ld f1, 0(r1)
       add r1, 8
       fadd f2, f0, f1
       sd f2, 0(r2)
       add r2, 8
       bne r1, r3, loop
```

Schedule

loop:

Int1	Int 2	M1	M2	FP+	FPx
add r1		ld			
				fadd	
add r2	bne	sd			

How many FP ops/cycle?

Loop Unrolling

Krste Asanovic
April 4, 2001
6.823, L14--11

```
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

Krste Asanovic
April 4, 2001
6.823, L14--12

Unroll 4 ways

```
loop:  ld f1, 0(r1)
       ld f2, 8(r1)
       ld f3, 16(r1)
       ld f4, 24(r1)
       add r1, 32
       fadd f5, f0, f1
       fadd f6, f0, f2
       fadd f7, f0, f3
       fadd f8, f0, f4
       sd f5, 0(r2)
       sd f6, 8(r2)
       sd f7, 16(r2)
       sd f8, 24(r2)
       add r2, 32
       bne r1, r3, loop
```

Schedule

	Int1	Int 2	M1	M2	FP+	FPx
loop:						

Scheduling Loop Unrolled Code

Unroll 4 ways

```
loop: ld f1, 0(r1)
      ld f2, 8(r1)
      ld f3, 16(r1)
      ld f4, 24(r1)
      add r1, 32
      fadd f5, f0, f1
      fadd f6, f0, f2
      fadd f7, f0, f3
      fadd f8, f0, f4
      sd f5, 0(r2)
      sd f6, 8(r2)
      sd f7, 16(r2)
      sd f8, 24(r2)
      add r2, 32
      bne r1, r3, loop
```

Schedule →

loop:

Int1	Int 2	M1	M2	FP+	FPx
		ld f1			
		ld f2			
		ld f3			
add r1		ld f4		fadd f5	
				fadd f6	
				fadd f7	
				fadd f8	
		sd f5			
		sd f6			
		sd f7			
add r2	bne	sd f8			

How many FLOPS/cycle?

Software Pipelining

Unroll 4 ways first

```
loop: ld f1, 0(r1)
      ld f2, 8(r1)
      ld f3, 16(r1)
      ld f4, 24(r1)
      add r1, 32
      fadd f5, f0, f1
      fadd f6, f0, f2
      fadd f7, f0, f3
      fadd f8, f0, f4
      sd f5, 0(r2)
      sd f6, 8(r2)
      sd f7, 16(r2)
      add r2, 32
      sd f8, -8(r2)
      bne r1, r3, loop
```

→

Int1	Int 2	M1	M2	FP+	FPx

Krste Asanovic
April 4, 2001
6.823, L14--15

Software Pipelining

Unroll 4 ways first

```
loop:  ld f1, 0(r1)
       ld f2, 8(r1)
       ld f3, 16(r1)
       ld f4, 24(r1)
       add r1, 32
       fadd f5, f0, f1
       fadd f6, f0, f2
       fadd f7, f0, f3
       fadd f8, f0, f4
       sd f5, 0(r2)
       sd f6, 8(r2)
       sd f7, 16(r2)
       add r2, 32
       sd f8, -8(r2)
       bne r1, r3, loop
```

	Int1	Int 2	M1	M2	FP+	FPx
prolog			ld f1			
			ld f2			
			ld f3			
	add r1		ld f4			
			ld f1		fadd f5	
			ld f2		fadd f6	
			ld f3		fadd f7	
	add r1		ld f4		fadd f8	
iterate, loop:			ld f1	sd f5	fadd f5	
			ld f2	sd f6	fadd f6	
		add r2	ld f3	sd f7	fadd f7	
	add r1	bne	ld f4	sd f8	fadd f8	
epilog				sd f5	fadd f5	
				sd f6	fadd f6	
		add r2		sd f7	fadd f7	
		bne		sd f8	fadd f8	
				sd f5		

How many FLOPS/cycle?

Krste Asanovic
April 4, 2001
6.823, L14--16

Software Pipelining vs. Loop Unrolling

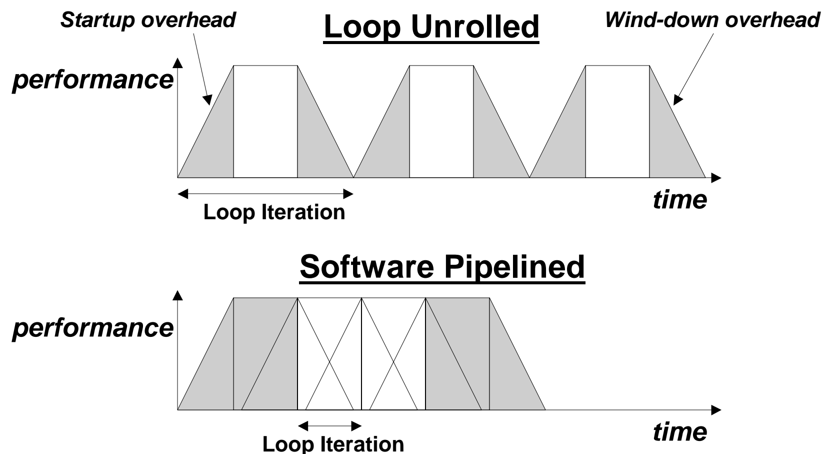

- ***Software pipelining pays startup/wind-down costs only once per loop, not once per iteration***

What if there are no loops?

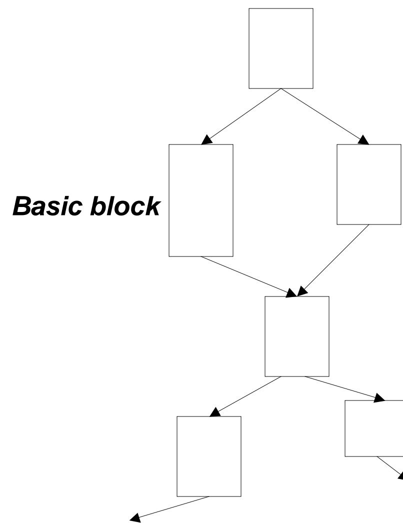

- Branches limit basic block size in control-flow intensive irregular code
- Difficult to find ILP in individual basic blocks

Trace Scheduling *[Fisher,Ellis]*

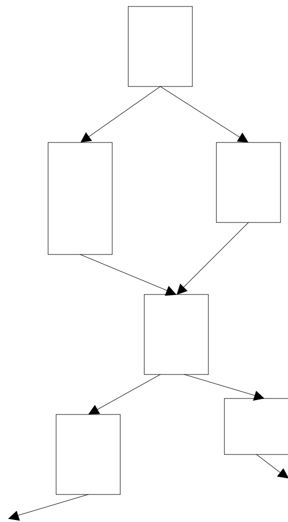

Trace Scheduling *[Fisher,Ellis]*

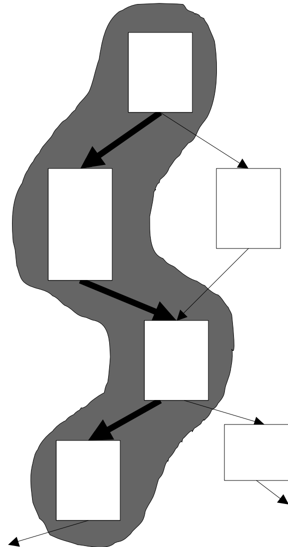

- Pick string of basic blocks, a *trace*, that represents most frequent branch path
- Use profiling feedback or compiler heuristics to find common branch paths
- Schedule whole "trace" at once
- Add fixup code to cope with branches jumping out of trace

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
- Scheduling around statically unpredictable branches
  - optimal schedule varies with branch path

VLIW Instruction Encoding

Various schemes to reduce effect of unused fields

- Compressed format in memory, expand on I-cache refill
- Cydra-5 MultiOp instructions: execute VLIW as sequential operations
- Mark parallel groups (used in TMS320C6x DSPs, Intel IA-64)

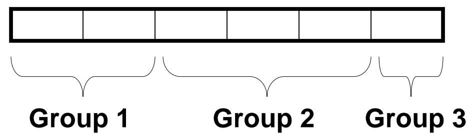

Rotating Register Files

Problem: Scheduled loops require lots of registers

Solution: Allocate new set of registers for each loop iteration

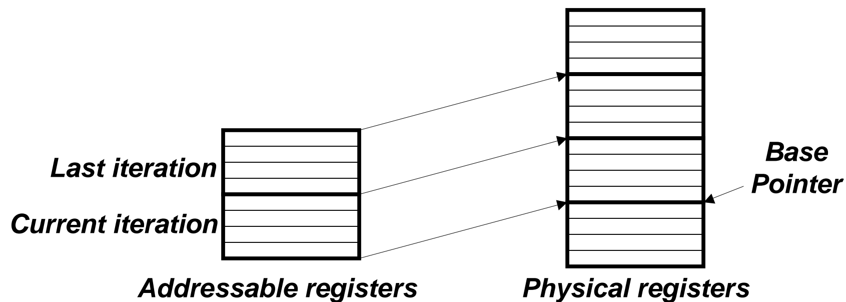

- Cheap form of register renaming
- Reduces code size in loop unrolled or software pipelined code
- Used in Cydra-5, Hitachi PA-RISC, ARM Piccolo DSP, Intel IA-64

Cydra-5: Memory Latency Register (MLR)

Problem: Loads have variable latency

Solution: Let software choose desired memory latency

- Compiler tries to schedule code for maximum load-use distance
- Software sets MLR to latency that matches code schedule
- Hardware ensures that loads take exactly MLR cycles to return values into processor pipeline
  - Hardware buffers loads that return early
  - Hardware stalls processor if loads return late

Intel EPIC IA-64

- EPIC is the style of architecture (cf. CISC, RISC)
  - Explicitly Parallel Instruction Computing
- IA-64 is Intel's chosen ISA (cf. x86, MIPS)
  - IA-64 = Intel Architecture 64-bit
  - An object-code compatible VLIW
- Itanium (aka Merced) is first implementation (cf. 8086)
  - First customer shipment should be in ~~1997~~ , ~~1998~~, ~~1999~~, ~~2000~~, 2001
  - McKinley will be second implementation due 2002

IA-64 Instruction Format

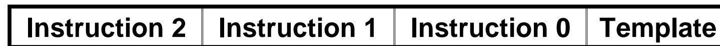

128-bit instruction bundle

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

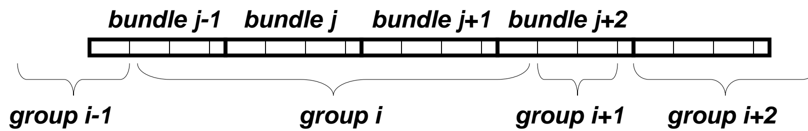

IA-64 Registers

- 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

Krste Asanovic
April 4, 2001
6.823, L14--27

IA-64 Predicated Execution

Problem: Mispredicted branches limit ILP

Solution: Eliminate some 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

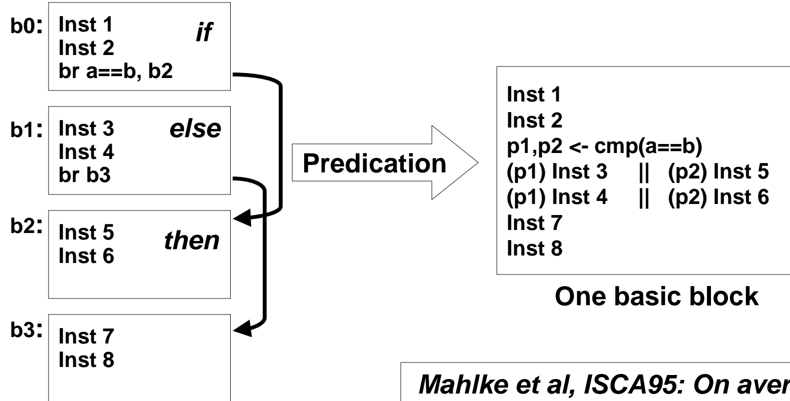

Mahlke et al, ISCA95: On average >50% branches removed

Krste Asanovic
April 4, 2001
6.823, L14--28

IA-64 Speculative Execution

Problem: Branches restrict compiler code motion

Solution: Speculative operations that don't cause exceptions

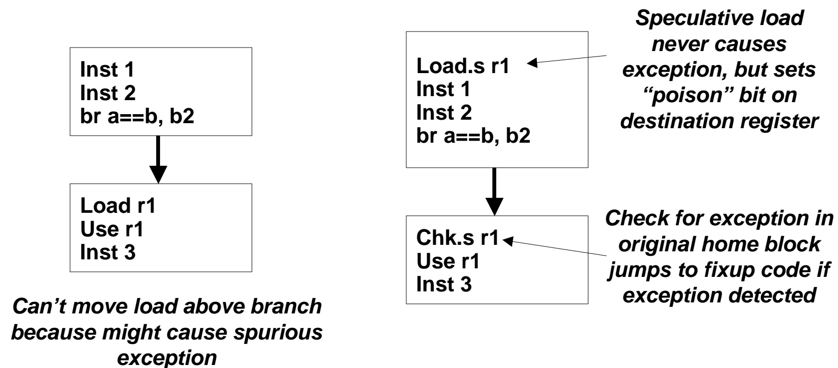

Particularly useful for scheduling long latency loads early

IA-64 Data Speculation

Problem: Possible memory hazards limit code scheduling

Solution: Hardware to check pointer hazards

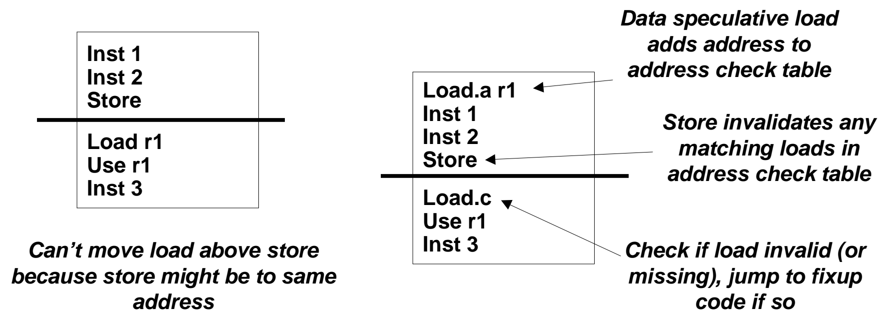

Requires associative hardware in address check table

ILP Datapath Hardware Scaling

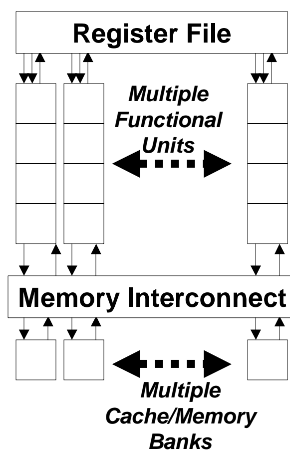

- **Replicating functional units and cache/memory banks is straightforward and scales linearly**
- **Register file ports and bypass logic for N functional units scale quadratically (N*N)**
- **Memory interconnection among N functional units and memory banks also scales quadratically**
- ***(For large N, could try O(N logN) interconnect schemes)***
- **Technology scaling: Wires are getting even slower relative to gate delays**
- **Complex interconnect adds latency as well as area**

=> *Need greater parallelism to hide latencies*

Clustered VLIW

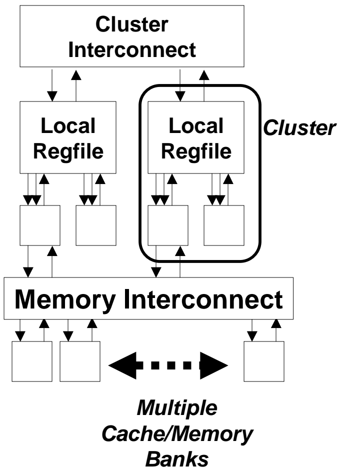

- Divide machine into clusters of local register files and local functional units
- Lower bandwidth/higher latency interconnect between clusters
- Software responsible for mapping computations to minimize communication overhead

Limits of Static Scheduling

Four major weaknesses of static scheduling:

- Unpredictable branches
- Variable memory latency (unpredictable cache misses)
- Code size explosion
- Compiler complexity