

# Computer Architecture


## Lecture 09 – VLIW (Instruction level Parallelism)

**Tian Xia**

Institute of Artificial Intelligence and Robotics  
Xi'an Jiaotong University

<http://gr.xjtu.edu.cn/web/pengjuren>


# Superscalar Control Logic Scaling



- Each issued instruction must somehow check against  $W \times L$  instructions to wake them up, i.e., growth in **hardware**  $\propto W \times (W \times L)$
- For **in-order** machines,  $L$  is related to pipeline latencies and check is done during issue (scoreboard)
- For **out-of-order** machines,  $L$  also includes time spent in IQ 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  $\Rightarrow$  greater  $L$   
 *$\Rightarrow$  Out-of-order control logic grows faster than  $W^2$  ( $\sim W^3$ )*

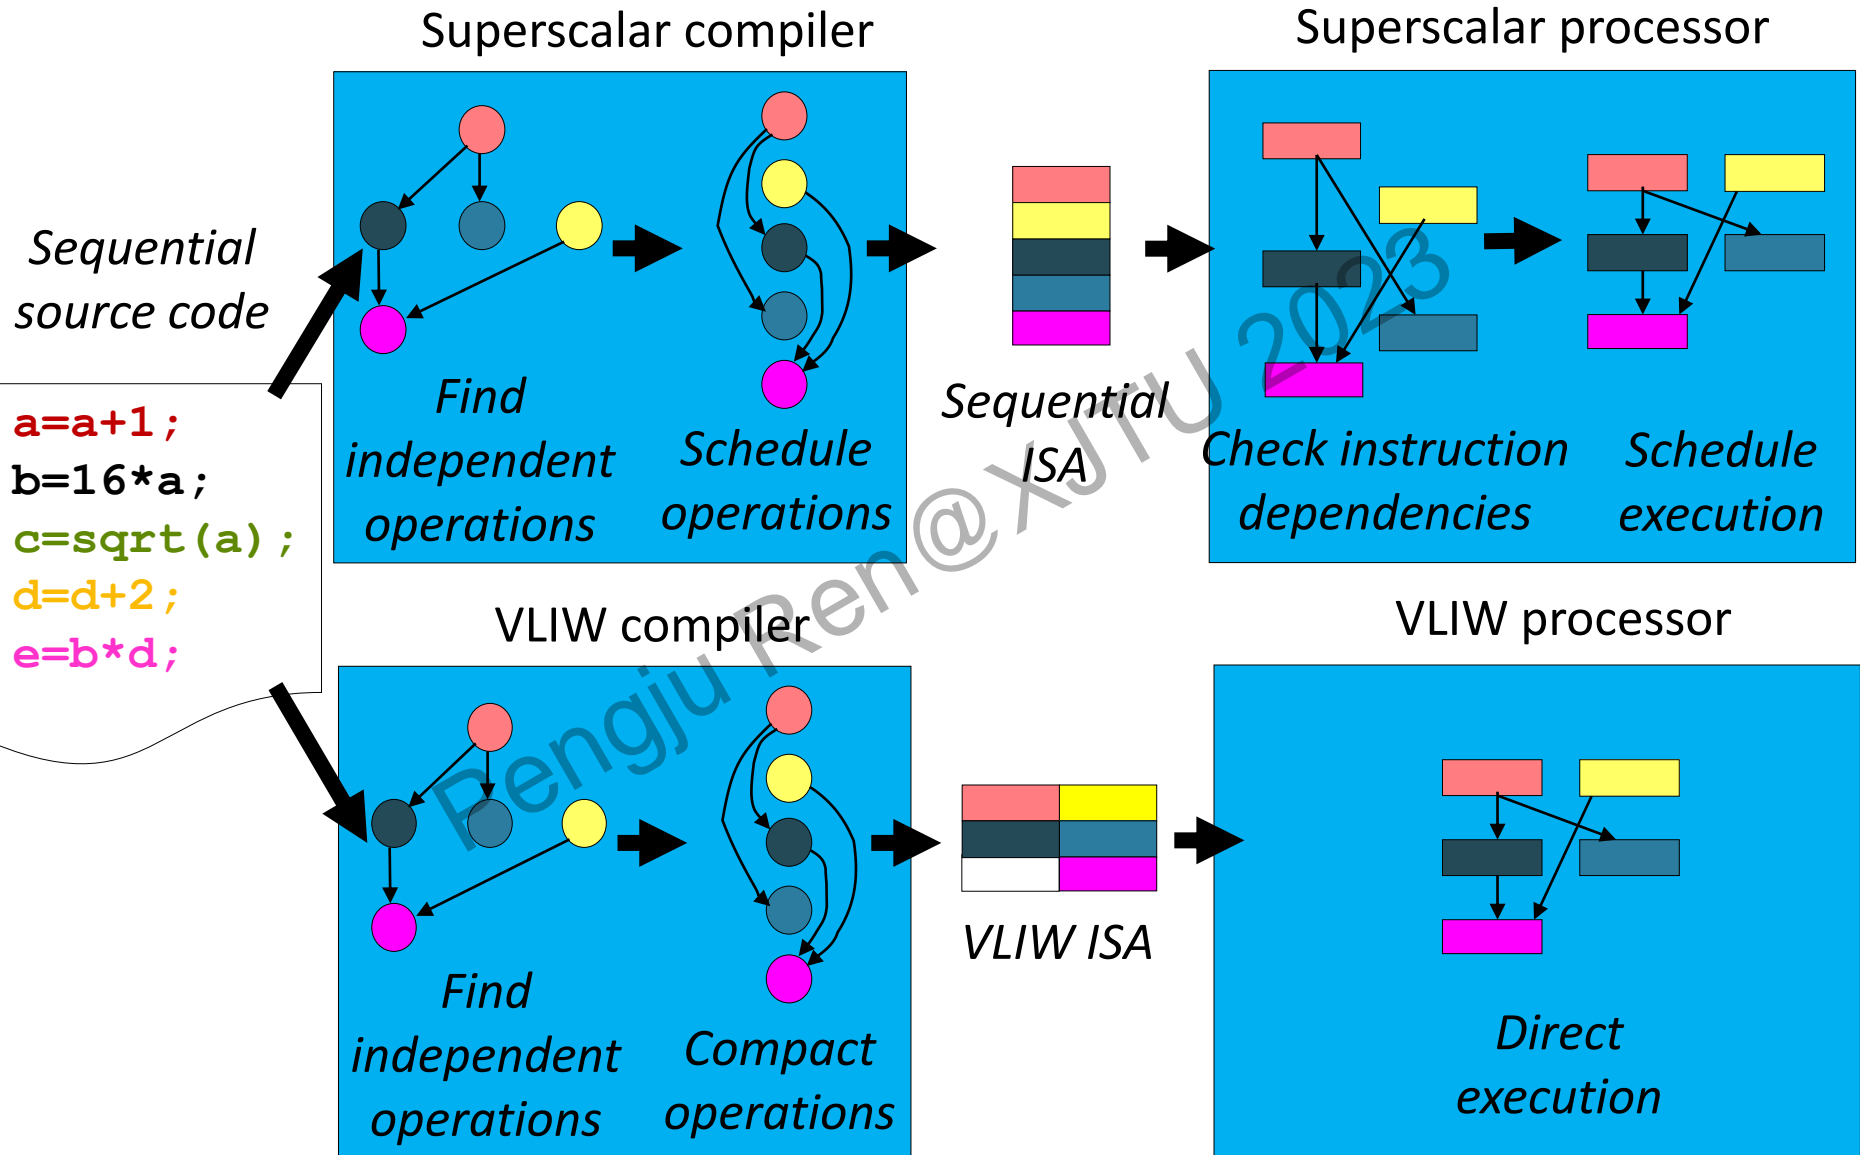
# Out-of-Order Control Complexity: MIPS R10000

*Control  
Logic*

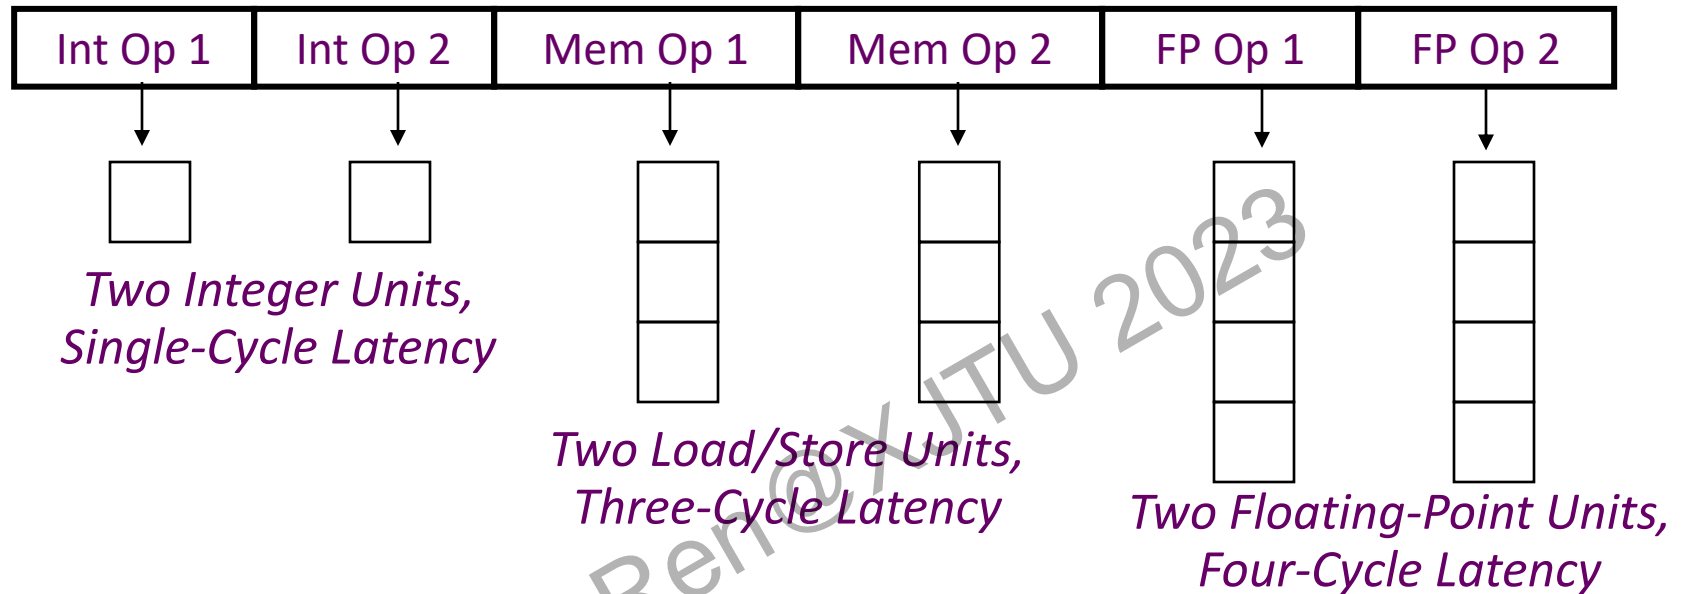


*[ SGI/MIPS Technologies  
Inc., 1995 ]*

# Sequential ISA Bottleneck



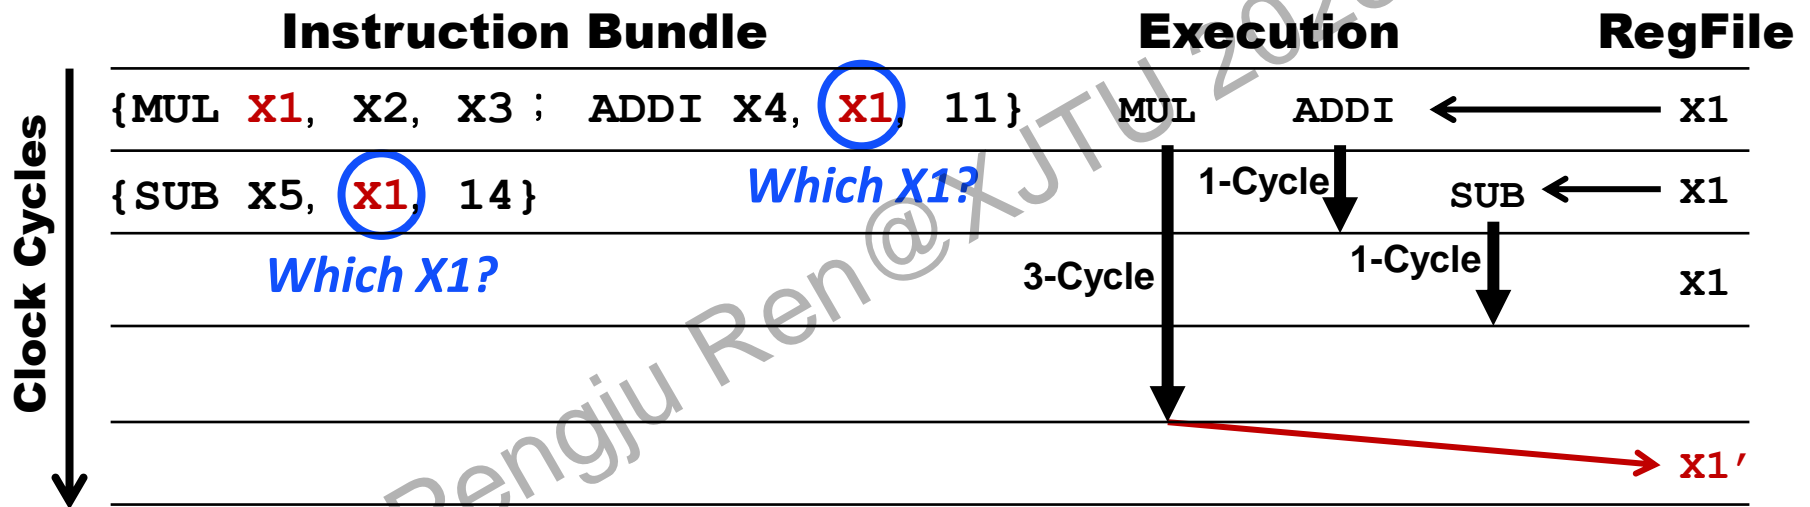
# VLIW: Very Long Instruction Word



- Multiple operations packed into one instruction (**bundle**)
- Each operation slot is for a **fixed function unit type**
- **Constant operation latencies** are specified (e.g. 3 for Mem Op)
- Architecture requires guarantee of:
  - Parallelism within an instruction => no cross-operation RAW check
  - No data use before data ready => no data interlocks

# VLIW: Very Long Instruction Word

- Architecture requires guarantee of:
  - Parallelism within an instruction => **no cross-operation RAW check**
  - No data use before data ready => **no data interlocks**




- Hardware **assumes compiler** is aware of all the **data dependency** and instruction **latency** of bundles

# VLIW Performance Example (2-wide bundles)

```
lw  $t0, 40($s0)    add $t1, $s1, $s2
sub $t2, $s1, $s3    and $t3, $s3, $s4
or  $t4, $s1, $s5    sw  $s5, 80($s0)
```

***Ideal IPC = 2***



***Actual IPC = 2*** (6 instructions issued in 3 cycles)



# VLIW Compiler Responsibilities

- **Statically schedule** (reorder) operations to maximize parallel execution
- Guarantees intra-instruction parallelism
- Schedule to avoid **data hazards** (no interlocks)
  - Typically separates operations with explicit NOPs
- What if data **cache miss** occurs?
  - (Scratchpad) Usually control cache manually by programmer
  - Could stop the pipeline and wait



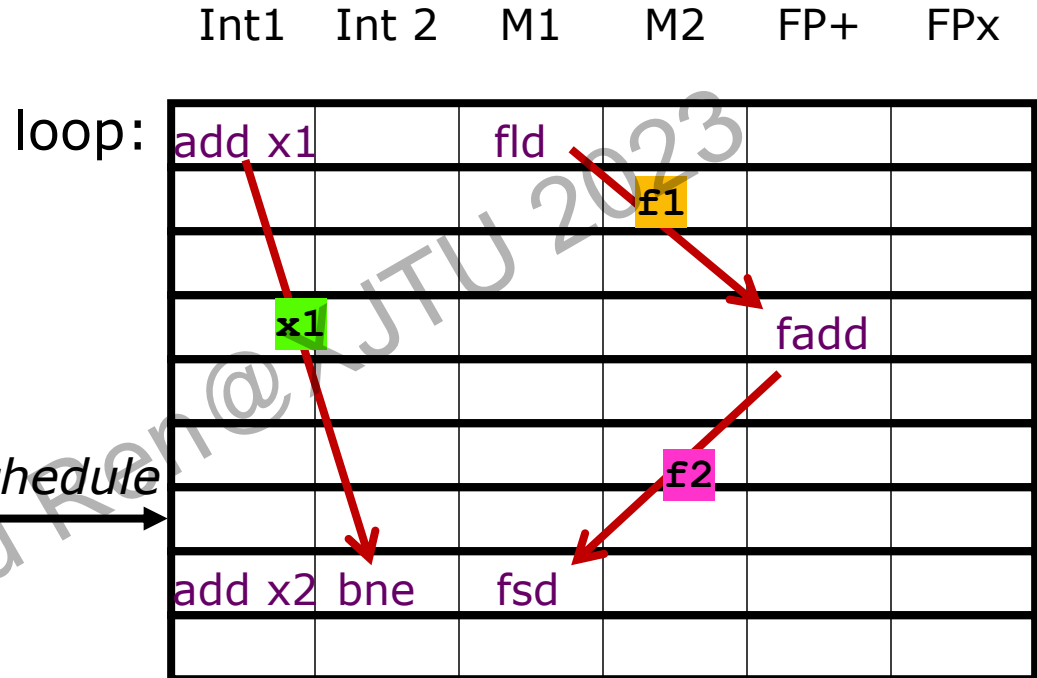
# Loop Execution

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

*Compile*

```
loop: fld f1, 0(x1)
      add x1, 8
      fadd f2, f0, f1
      fsd f2, 0(x2)
      add x2, 8
      bne x1, x3, loop
```

*Schedule*



How many FP ops/cycle?

$$1 \text{ fadd} / 8 \text{ cycles} = 0.125$$

# 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;  
}
```

*Is this code correct ?*

- Need to handle values of N that are not multiples of unrolling factor with **final cleanup** loop
- More instructions (→ larger I cache)

# Scheduling Loop Unrolled Code

**Unroll into 4 ways**

```

loop: fld f1, 0(x1)
      fld f2, 8(x1)
      fld f3, 16(x1)
      fld f4, 24(x1)
      add x1, 32
      fadd f5, f0, f1
      fadd f6, f0, f2
      fadd f7, f0, f3
      fadd f8, f0, f4
      fsd f5, 0(x2)
      fsd f6, 8(x2)
      fsd f7, 16(x2)
      fsd f8, 24(x2)
      add x2, 32
      bne x1, x3, loop
    
```

loop:

*Schedule* →

Int1	Int 2	M1	M2	FP+	FPx
		fld f1			
		fld f2			
		fld f3			
add x1		fld f4		fadd f5	
				fadd f6	
				fadd f7	
				fadd f8	
		fsd f5			
		fsd f6			
		fsd f7			
add x2	bne	fsd f8			

How many FLOPS/cycle?


$$4 \text{ fadds} / 11 \text{ cycles} = 0.36$$

# Software Pipelining

Exploit independent loop iterations


- If loop iterations are **independent**, then get more parallelism by scheduling instructions from **different iterations**
- Construct the data-flow graph for one iteration

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

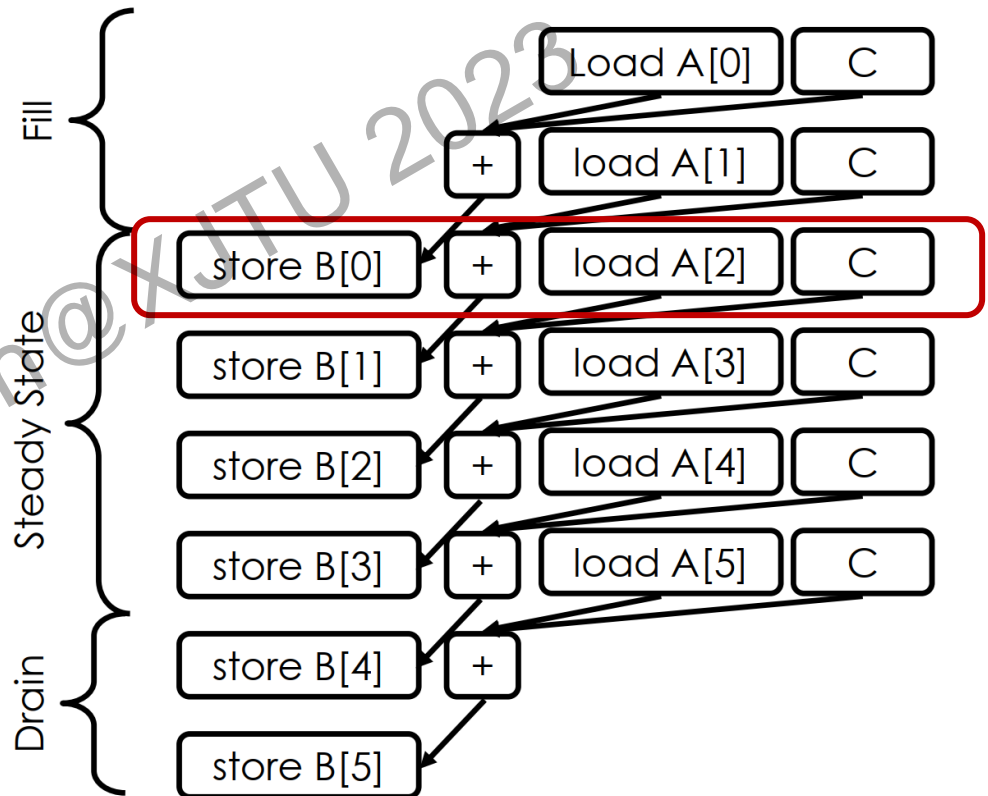


# Software Pipelining

Not pipelined



Pipelined



# Software Pipelining

*Unroll 4 ways first*

```

loop: fld f1, 0(x1)
      fld f2, 8(x1)
      fld f3, 16(x1)
      fld f4, 24(x1)
      add x1, 32
      fadd f5, f0, f1
      fadd f6, f0, f2
      fadd f7, f0, f3
      fadd f8, f0, f4
      fsd f5, 0(x2)
      fsd f6, 8(x2)
      fsd f7, 16(x2)
      fsd f8, 24(x2)
      add x2, 32
      bne x1, x3, loop
    
```

**loop:**

add x1

Conflict: put  
BNE in the  
next bundle

Int1    Int 2    M1    M2    FP+    FPx

		fld f1			
		fld f2			
		fld f3			
add x1		fld f4			
		fld f1		fadd f5	
		fld f2		fadd f6	
		fld f3		fadd f7	
add x1		fld f4		fadd f8	
		fld f1	fsd f5	fadd f5	
		fld f2	fsd f6	fadd f6	
		fld f3	fsd f7	fadd f7	
add x1	add x2 bne	fld f4	fsd f8	fadd f8	
			fsd f5	fadd f5	
			fsd f6	fadd f6	
			fsd f7	fadd f7	
add x2 bne			fsd f8	fadd f8	
			fsd f5		

# Software Pipelining

*Unroll 4 ways first*

```

loop: fld f1, 0(x1)
      fld f2, 8(x1)
      fld f3, 16(x1)
      fld f4, 24(x1)
      add x1, 32
      fadd f5, f0, f1
      fadd f6, f0, f2
      fadd f7, f0, f3
      fadd f8, f0, f4
      fsd f5, 0(x2)
      fsd f6, 8(x2)
      fsd f7, 16(x2)
      add x2, 32
      fsd f8, -8(x2)
      bne x1, x3, loop
    
```

**loop:**

Int1	Int 2	M1	M2	FP+	FPx
		fld f1			
		fld f2			
		fld f3			
add x1		fld f4			
		fld f1		fadd f5	
		fld f2		fadd f6	
		fld f3		fadd f7	
add x1		fld f4		fadd f8	
		fld f1	fsd f5	fadd f5	
		fld f2	fsd f6	fadd f6	
		add x2	fsd f7	fadd f7	
add x1	bne	fld f4	fsd f8	fadd f8	
			fsd f5	fadd f5	
			fsd f6	fadd f6	
	add x2		fsd f7	fadd f7	
	bne		fsd f8	fadd f8	
			fsd f5		

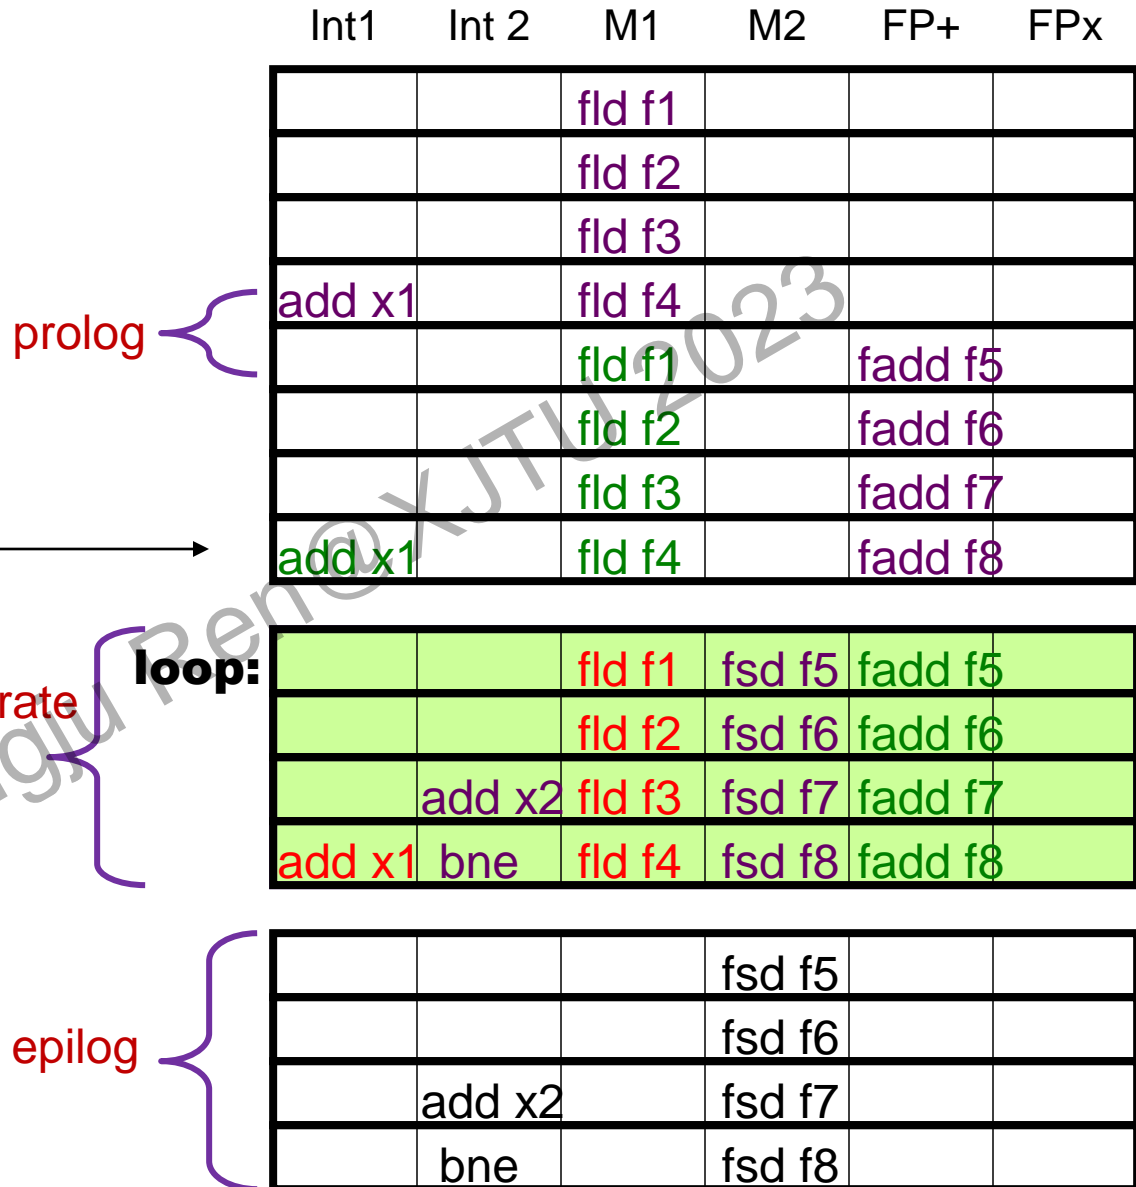


# Software Pipelining

*Unroll 4 ways first*


```

loop: fld f1, 0(x1)
      fld f2, 8(x1)
      fld f3, 16(x1)
      fld f4, 24(x1)
      add x1, 32
      fadd f5, f0, f1
      fadd f6, f0, f2
      fadd f7, f0, f3
      fadd f8, f0, f4
      fsd f5, 0(x2)
      fsd f6, 8(x2)
      fsd f7, 16(x2)
      add x2, 32
      fsd f8, -8(x2)
      bne x1, x3, loop
    
```




# Software Pipelining vs. Loop Unrolling

## Loop Unrolled




## Software Pipelined




*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]

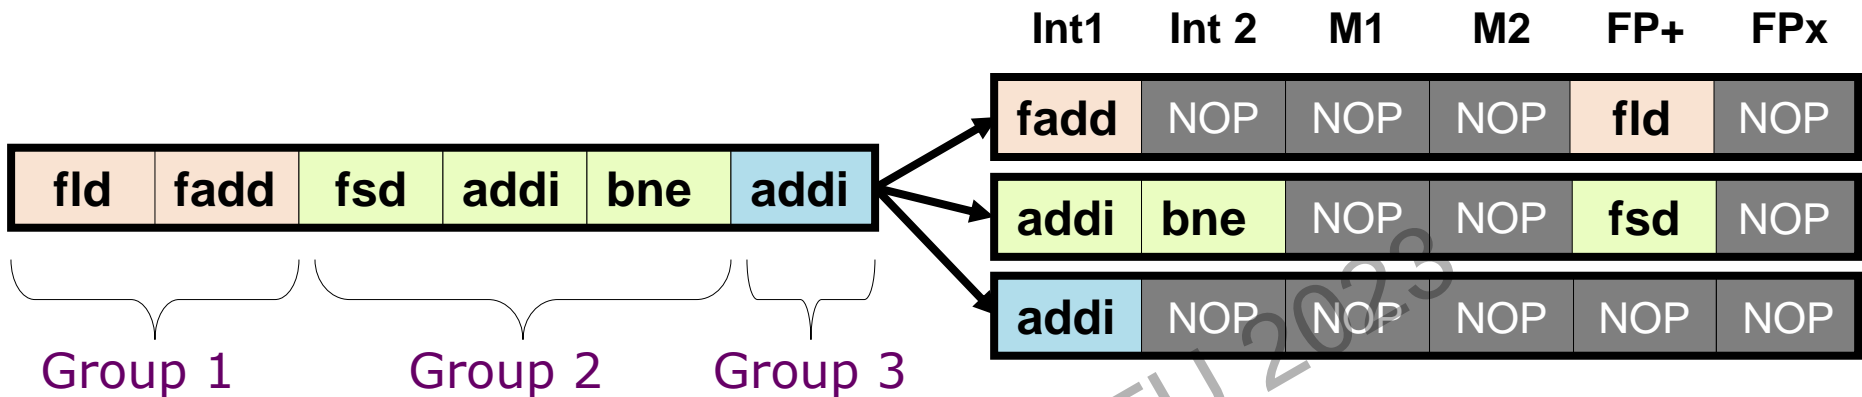


- A trace is a **possible sequence** of **basic blocks** (a.k.a., long string of straight-line code)
- Trace Selection: Use **profiling** or **compiler heuristics** to find common sequences/paths
- **Trace Compaction**: Schedule whole trace into few VLIW instructions
- Add **fixup code** to cope with 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
- Knowing branch probabilities
  - **Profiling** requires an significant extra step in build process
- Scheduling for statically **unpredictable** branches
  - optimal schedule varies with branch path
- Object code size
  - **instruction padding (NOPs)** 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
  - Uncertainty about addresses limit code reordering

# VLIW Instruction Encoding



- Schemes to reduce effect of unused fields (NOPs)
  - **Compressed format** in memory, expand (**uncompress**) 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
- 544mm<sup>2</sup> 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



128-bit instruction bundle

- Template bits describe grouping of these instructions with others in adjacent bundles



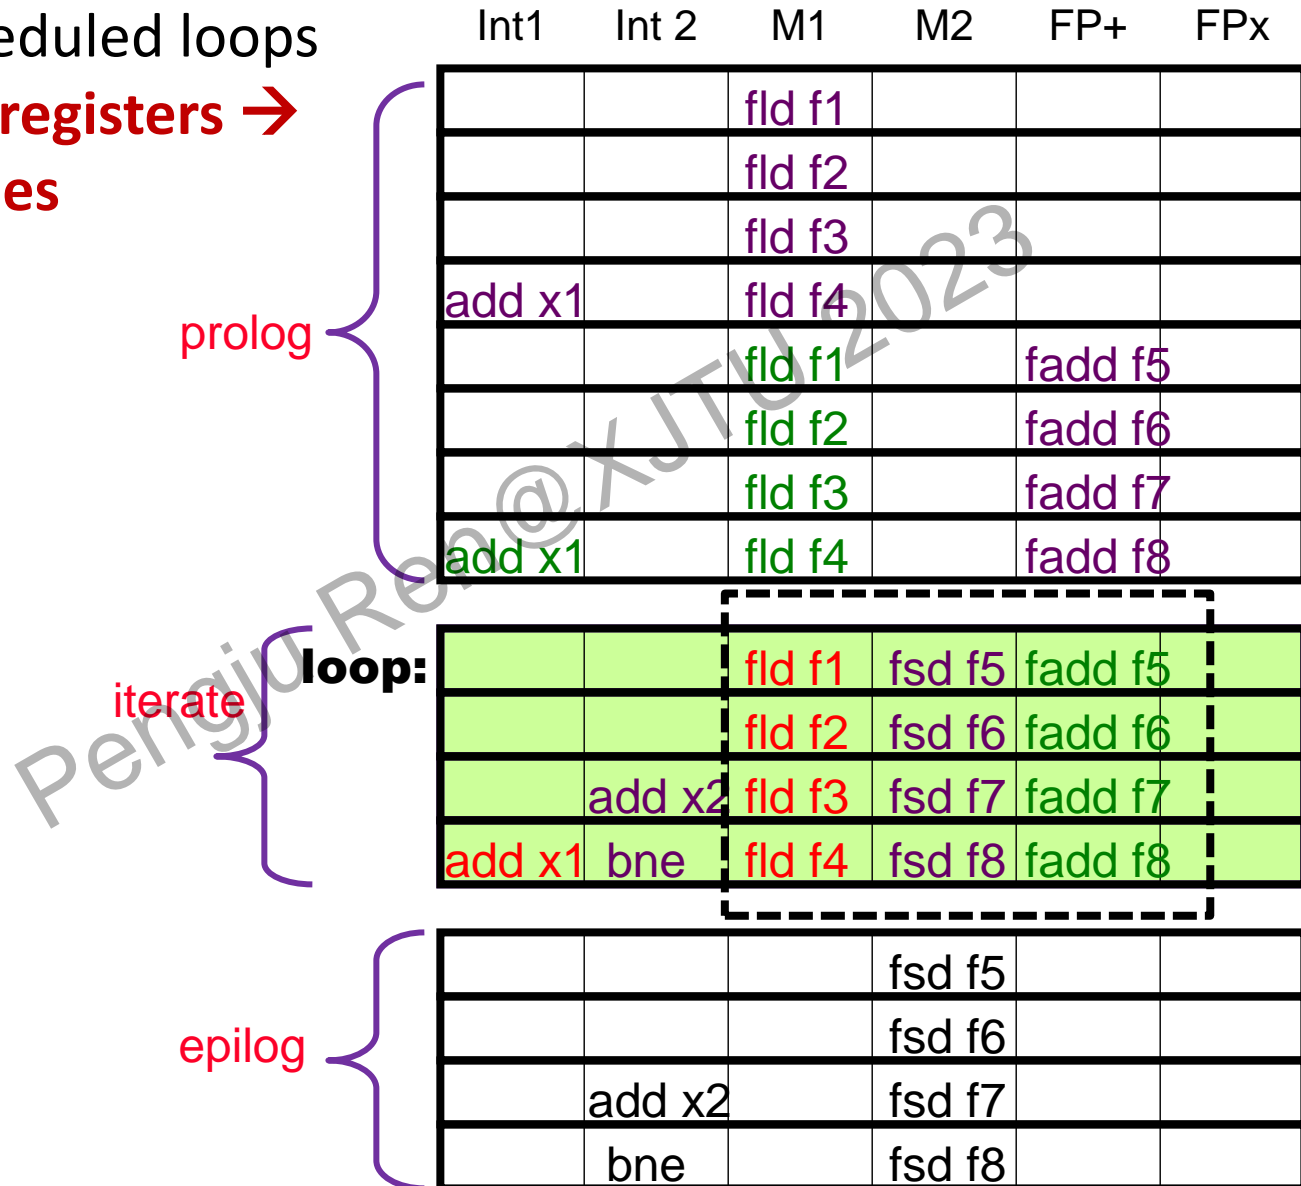
- Each group contains instructions that can execute in parallel

# 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
  - Rotation is a simple form of **register renaming** allowing one instruction to address different physical registers on each iteration

# Recap: Soft-Pipelining


**Problems:** Scheduled loops  
require **lots of registers** →  
**duplicated codes**



# Rotating Register Files

**Problems:** Scheduled loops require **lots of registers** → **duplicated codes**

**Solution:** Automatically use new set of registers for each loop iteration



- **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 **non-rotating** registers.

# Rotating Register Files

## (Software/Hardware Co-design)

Register rotation is used for optimizing loops that are both counted or data-terminated.

- **Counted loops** are loops whose iterations are known prior to entering the loop
- **Data-terminated loops** are dependent upon values calculated inside the loop.

Register Set	Static	Rotating
General Registers (GR)	0-31	32-127
Floating Point Registers (FR)	0-31	32-127
Predicate Registers (PR)	0-15	16-63

The general, floating point and predicate registers are divided into subsets of **static** and **rotating** sets. The above is the subdivision

# Recap: Soft-Pipelining

**Problems:** Scheduled loops require **lots of registers** → **duplicated codes**

prolog

iterate

loop:

**Solution:** Reduce loop codes to **one instruction**

epilog

Int1    Int 2    M1    M2    FP+    FPx

		fld f1			
		fld f2			
		fld f3			
add x1		fld f4			
		fld f1		fadd f5	
		fld f2		fadd f6	
		fld f3		fadd f7	
add x1		fld f4		fadd f8	

		fld f1	fsd f5	fadd f5	
		fld f2	fsd f6	fadd f6	
	add x2	fld f3	fsd f7	fadd f7	
add x1 bne		fld f4	fsd f8	fadd f8	

Static

Rotating

			fsd f5		
			fsd f6		
	add x2		fsd f7		
	bne		fsd f8		



# 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 ( $f9 - f5 = 4$ )

**3 cycles later, f1 is renamed as f4**

**4 cycles later, f5 is renamed as f9**



RRB=8	<code>fld P9, ()</code>	<code>fadd P13, P12,</code>	<code>fsd P17, ()</code>	<code>bloop</code>
RRB=7	<code>fld P8, ()</code>	<code>fadd P12, P11,</code>	<code>fsd P16, ()</code>	<code>bloop</code>
RRB=6	<code>fld P7, ()</code>	<code>fadd P11, P10,</code>	<code>fsd P15, ()</code>	<code>bloop</code>
RRB=5	<code>fld P6, ()</code>	<code>fadd P10, P9,</code>	<code>fsd P14, ()</code>	<code>bloop</code>
RRB=4	<code>fld P5, ()</code>	<code>fadd P9, P8,</code>	<code>fsd P13, ()</code>	<code>bloop</code>
RRB=3	<code>fld P4, ()</code>	<code>fadd P8, P7,</code>	<code>fsd P12, ()</code>	<code>bloop</code>
RRB=2	<code>fld P3, ()</code>	<code>fadd P7, P6,</code>	<code>fsd P11, ()</code>	<code>bloop</code>
RRB=1	<code>fld P2, ()</code>	<code>fadd P6, P5,</code>	<code>fsd P10, ()</code>	<code>bloop</code>

Annotations and Renaming:

- F1=P9** (at RRB=8)
- F5=P10** (at RRB=5)
- F4=P9** (at RRB=5)
- F9=P10** (at RRB=1)

Only need **1 instruction** in loop code!

# IA-64 Predicated Execution

- **Predication** is the **conditional execution** of instructions.
- In **traditional** architectures, conditional execution is implemented through **branches**.
- In **VLIW** machine, **predicated execution** avoids branches, and simplifies compiler optimization by converting a **control dependency** to a **data dependency**.

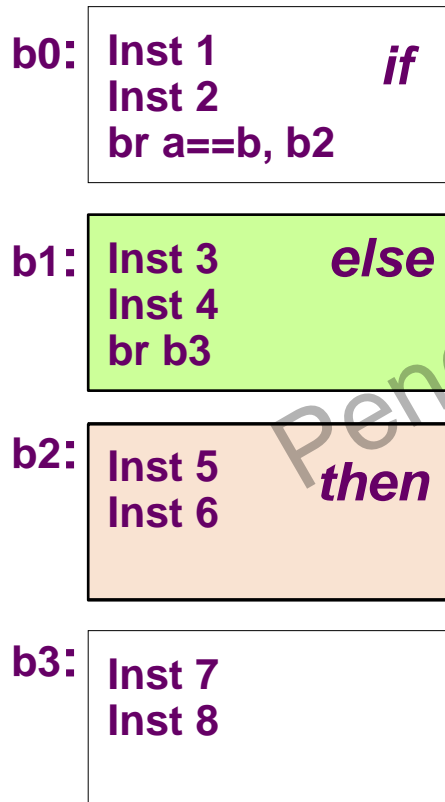
# IA-64 Predicated Execution

**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

4 basic blocks



1 basic block

Predication



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


# IA-64 Normal Compares

**Compare operation** works on a **pair of predicate registers**.

- Compare operations play a key role in IA-64, and particularly in relation to predication.

**Normal Compare instruction** evaluates the expression and then:

- Set **the first predicate register** to the result of the comparison
- Set **the second predicate register** to the complement of the comparison.



```
If (a == b) {  
    c++;  
} else {  
    d++;  
}
```


Predication

```
{  
    cmp.eq p1, p2 = ra, rb  
    (p1) add rc = rc, 1  
    (p2) add rd = rd, 1  
}
```

# IA-64 Unconditional Compares

**Unconditional Compare** instructions are **predicated themselves**:

- When its **self-predicate is 1**, the compare executes normally and writes to its target registers as would a normal compare.
- When its **self-predicate is 0**, **write 0 to both** of its target registers.



Unconditional Compare is useful in **nested if-conversion**.

```

If (a > b) {
    c++;
} else {
    d+=c;
    if (e==f) {
        g++;
    } else {
        h--;
    }
}
    
```

Predication

```

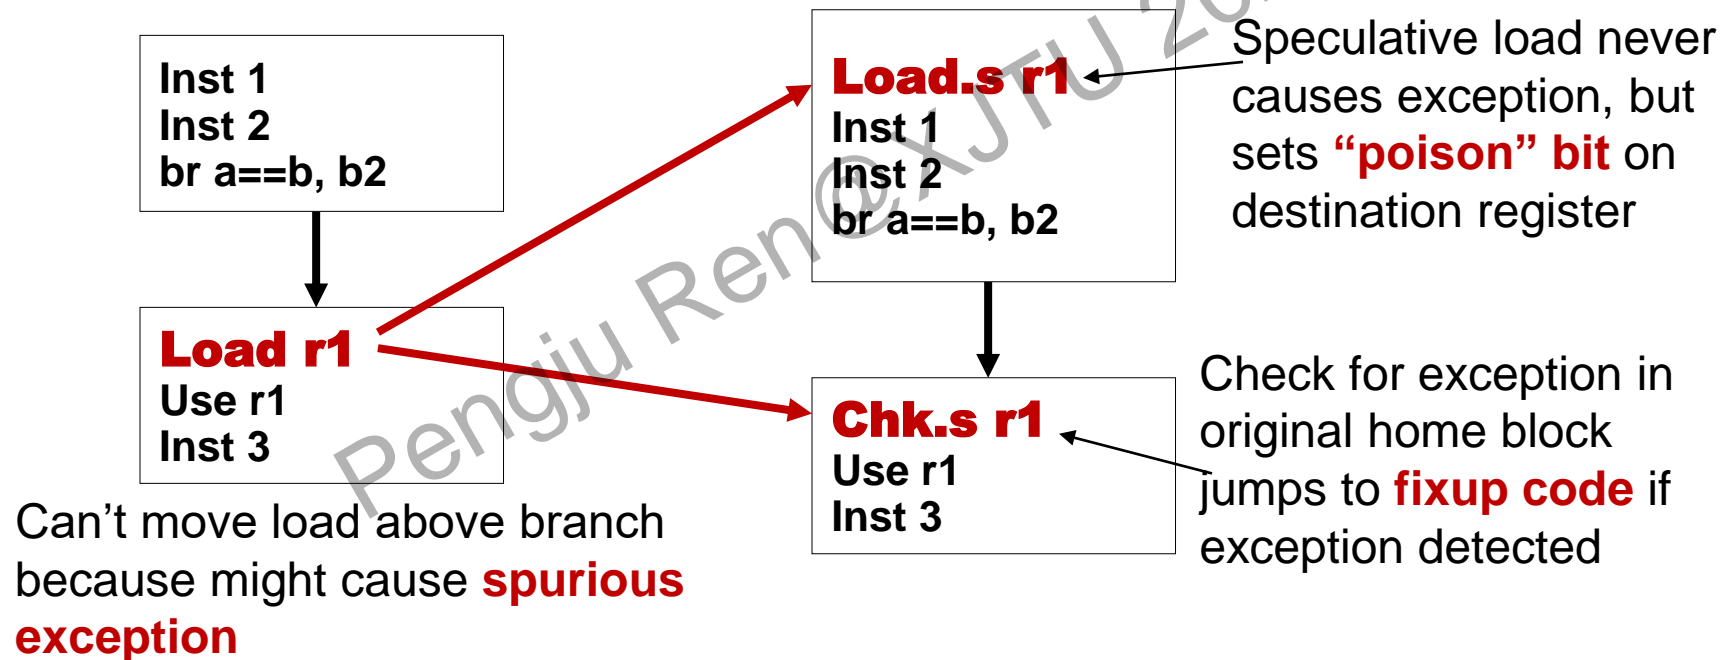
{
    cmp.gt p1, p2 = ra, rb
} {
    (p1) add rc = rc, 1
    (p2) add rd = rd, rc
    (p2) cmp.eq.unc p3, p4 = re, rf
} {
    (p3) add rg = rg, 1
    (p4) add rh = rh, -1
}
    
```

# IA-64 Speculative Execution

**Problem:** Branches restrict compiler code motion

**Solution:** Speculative operations that don't cause exceptions

- Requires associative hardware in register poison bit
- Particularly useful for scheduling long latency loads early



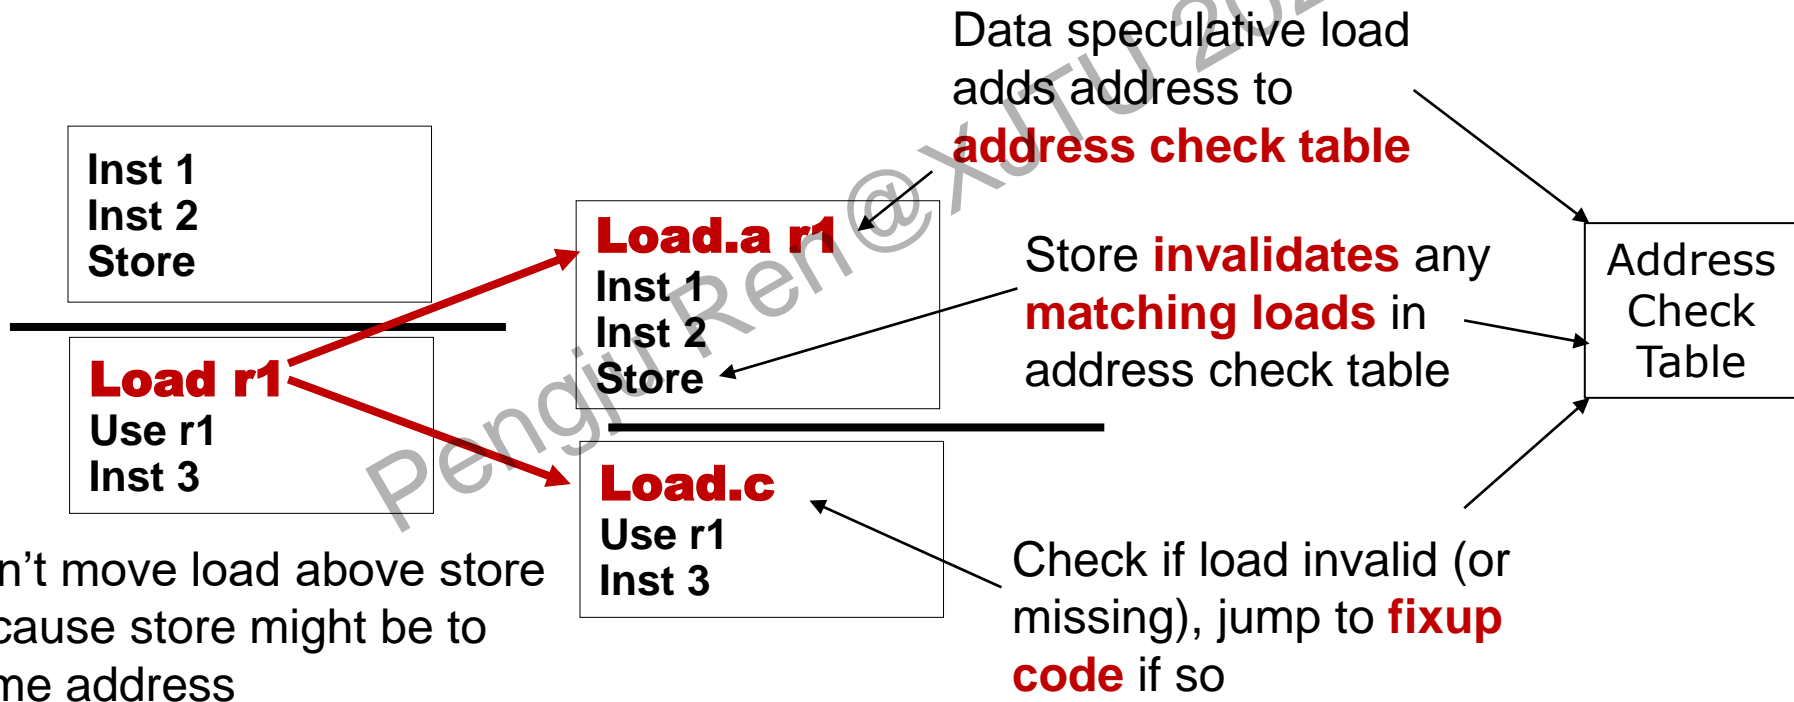
- Compiler guarantees the **next reference of r1 is not as src register** except for check
- If branch is not taken and check is never executed, **clear poison bit** when next time r1 is modified.

# IA-64 Data Speculation

**Problem:** Possible memory hazards limit code scheduling

**Solution:** Hardware to check pointer hazards

-- Requires associative hardware in **address check table**





# Limits of Static Scheduling

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

*Next Lecture : Vectors and SIMD*  
*(Data Level Parallel)*