

CS252  
Graduate Computer Architecture  
Lecture 8

Instruction Level Parallelism 2:  
Getting the CPI < 1

September 22, 1999  
Prof. John Kubiawicz

9/22/99

CS252/Kubiawicz  
Lec 8.1

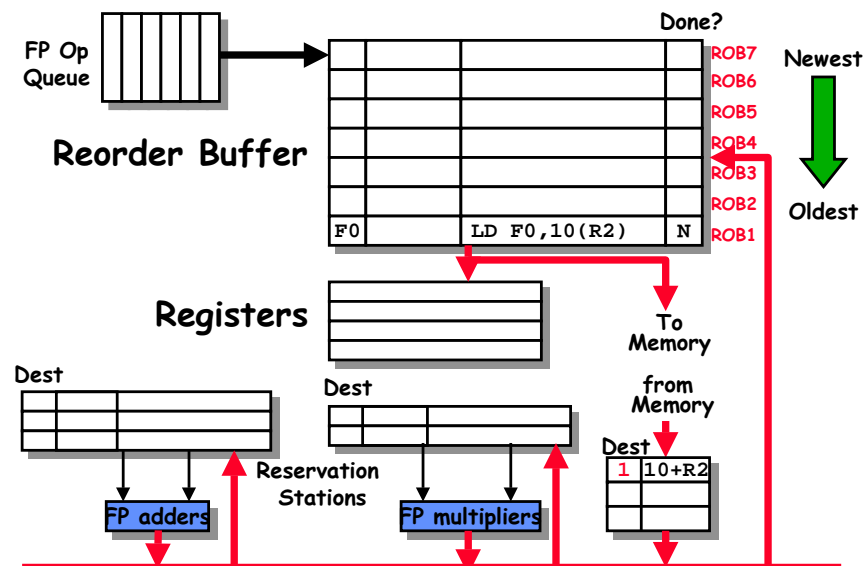
Review:  
Hardware unrolling, in-order commit,  
and explicit register renaming

- Machines that use hardware techniques with register renaming (such as tomasulo) can unroll loops automatically in hardware
- **In-Order-Commit** is important because:
  - Allows the generation of precise exceptions
  - Allows speculation across branches
- Use of reorder buffer
  - Commits user-visible state in instruction order
- Explicit register renaming uses a rename table and large bank of physical registers

9/22/99

CS252/Kubiawicz  
Lec 8.2

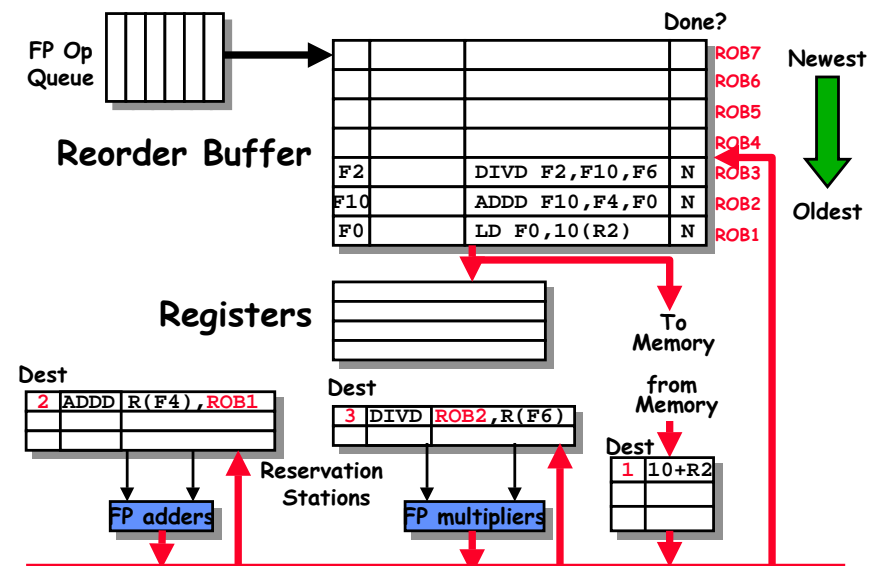
Tomasulo With Reorder buffer:



9/22/99

CS252/Kubiawicz  
Lec 8.3

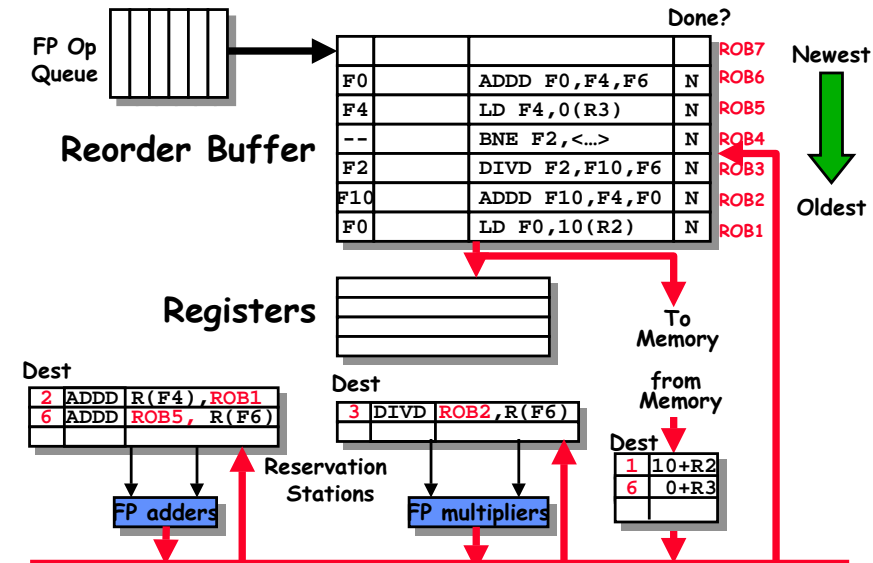
Tomasulo With Reorder buffer:



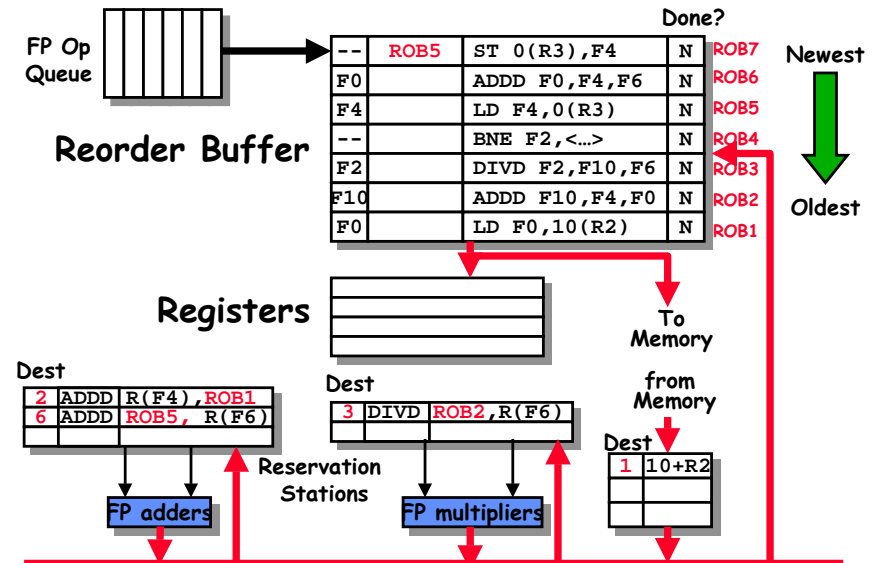
9/22/99

CS252/Kubiawicz  
Lec 8.4

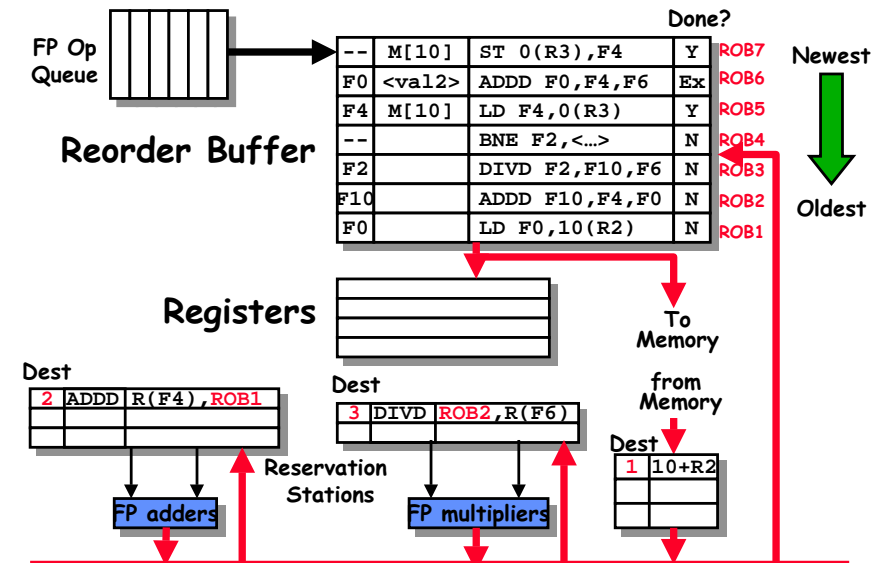
## Tomasulo With Reorder buffer:



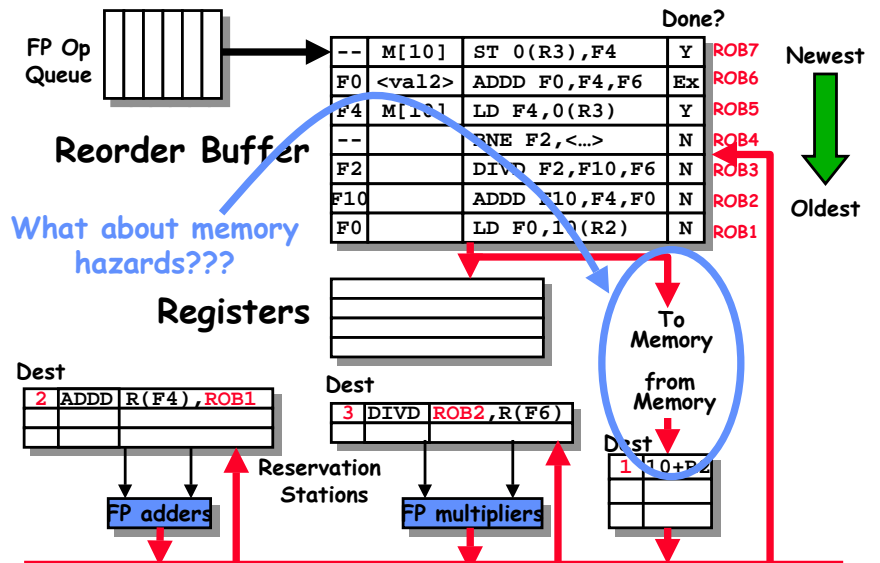
## Tomasulo With Reorder buffer:



## Tomasulo With Reorder buffer:



## Tomasulo With Reorder buffer:



## Memory Disambiguation: Sorting out RAW Hazards in memory

- **Question: Given a load that follows a store in program order, are the two related?**
  - (Alternatively: is there a RAW hazard between the store and the load)?

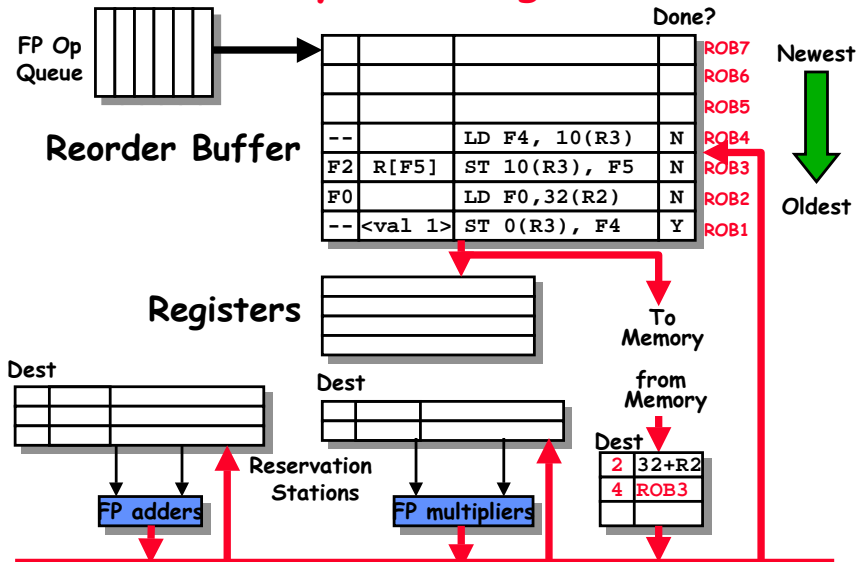
Eg:     st     0(R2),R5  
         ld     R6,0(R3)

- Can we go ahead and start the load early?
  - Store address could be delayed for a long time by some calculation that leads to R2 (divide?).
  - We might want to issue/begin execution of both operations in same cycle.
  - Today: Answer is that we are not allowed to start load until we know that address  $O(R2) \neq O(R3)$
  - Next Week: We might guess at whether or not they are dependent (called "dependence speculation") and use reorder buffer to fixup if we are wrong.

## Hardware Support for Memory Disambiguation


- Need buffer to keep track of all outstanding stores to memory, in program order.
  - Keep track of address (when becomes available) and value (when becomes available)
  - FIFO ordering: will retire stores from this buffer in program order
- When issuing a load, record current head of store queue (know which stores are ahead of you).
- When have address for load, check store queue:
  - If *any* store prior to load is waiting for its address, stall load.
  - If load address matches earlier store address (associative lookup), then we have a *memory-induced RAW hazard*:
    - » store value available  $\Rightarrow$  return value
    - » store value not available  $\Rightarrow$  return ROB number of source
  - Otherwise, send out request to memory
- Actual stores commit in order, so no worry about WAR/WAW hazards through memory.

## Memory Disambiguation:



## Explicit register renaming:

## Hardware equivalent of static, single-assignment (SSA) compiler form

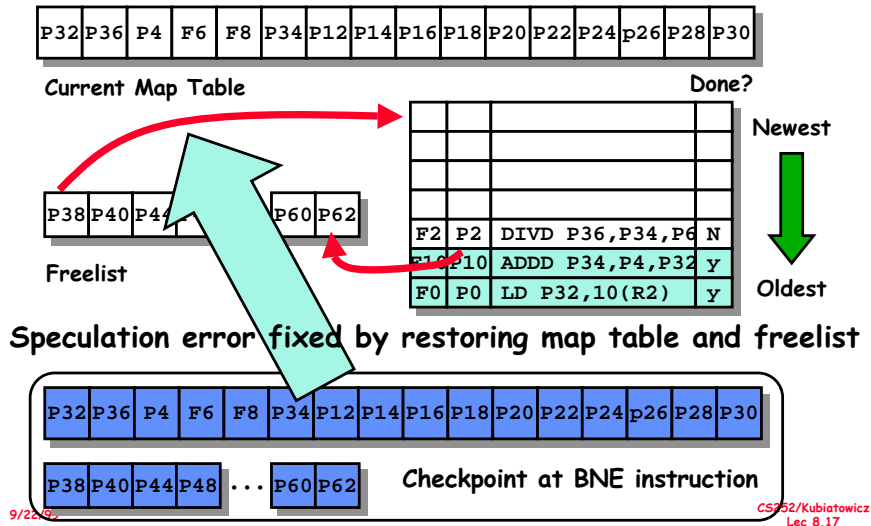


- Physical register file larger than ISA register file
- On issue, each instruction that modifies a register is allocated new physical register from freelist
- Used on: R10000, Alpha 21264, HP PA8000



## Explicit register renaming:

Hardware equivalent of static,  
single-assignment (SSA) compiler form



## Instruction Level Parallelism

- High speed execution based on *instruction level parallelism* (ilp): potential of short instruction sequences to execute in parallel
- High-speed microprocessors exploit ILP by:
  - 1) pipelined execution: overlap instructions
  - 2) Out-of-order execution (commit in-order)
  - 3) Multiple issue: issue and execute multiple instructions per clock cycle
  - 4) Vector instructions: many independent ops specified with a single instruction
- Memory accesses for high-speed microprocessor?
  - Data Cache possibly multiported, multiple levels

9/22/99

CS252/Kubiatowicz  
Lec 8.18

## CS 252 Administrivia

- Exam: Wednesday 10/13  
Location: TBA  
TIME: 5:30 - 8:30
- This info is on the Lecture page (has been)
- Meet at LaVal's afterwards for Pizza and Beverages
- Assignment next time
- Computers in the News: Intel IXP 2000 initiative

9/22/99

CS252/Kubiatowicz  
Lec 8.19

## Getting CPI < 1: Issuing Multiple Instructions/Cycle

- Two variations
- Superscalar: varying no. instructions/cycle (1 to 8), scheduled by compiler or by HW (Tomasulo)
  - IBM PowerPC, Sun UltraSparc, DEC Alpha, HP 8000
- (Very) Long Instruction Words (V)LIW: fixed number of instructions (4-16) scheduled by the compiler; put ops into wide templates
  - Joint HP/Intel agreement in 1999/2000?
  - Intel Architecture-64 (IA-64) 64-bit address
  - Style: "Explicitly Parallel Instruction Computer (EPIC)"
- Anticipated success lead to use of Instructions Per Clock cycle (IPC) vs. CPI

9/22/99

CS252/Kubiatowicz  
Lec 8.20

## Getting CPI < 1: Issuing Multiple Instructions/Cycle

- Superscalar DLX: 2 instructions, 1 FP & 1 anything else

- Fetch 64-bits/clock cycle; Int on left, FP on right
- Can only issue 2nd instruction if 1st instruction issues
- More ports for FP registers to do FP load & FP op in a pair

| Type             | Pipe Stages |    |    |     |     |     |    |
|------------------|-------------|----|----|-----|-----|-----|----|
| Int. instruction | IF          | ID | EX | MEM | WB  |     |    |
| FP instruction   | IF          | ID | EX | MEM | WB  |     |    |
| Int. instruction |             | IF | ID | EX  | MEM | WB  |    |
| FP instruction   |             | IF | ID | EX  | MEM | WB  |    |
| Int. instruction |             |    | IF | ID  | EX  | MEM | WB |
| FP instruction   |             |    | IF | ID  | EX  | MEM | WB |

- 1 cycle load delay expands to 3 instructions in SS
- instruction in right half can't use it, nor instructions in next slot

9/22/99

CS252/Kubiatowicz  
Lec 8.21

## Review: Unrolled Loop that Minimizes Stalls for Scalar

```

1 Loop: LD      F0,0(R1)           LD to ADDD: 1 Cycle
2         LD      F6,-8(R1)        ADDD to SD: 2 Cycles
3         LD      F10,-16(R1)
4         LD      F14,-24(R1)
5         ADDD    F4,F0,F2
6         ADDD    F8,F6,F2
7         ADDD    F12,F10,F2
8         ADDD    F16,F14,F2
9         SD      0(R1),F4
10        SD      -8(R1),F8
11        SD      -16(R1),F12
12        SUBI    R1,R1,#32
13        BNEZ    R1,LOOP
14        SD      8(R1),F16      ; 8-32 = -24
    
```

14 clock cycles, or 3.5 per iteration

9/22/99

CS252/Kubiatowicz  
Lec 8.22

## Loop Unrolling in Superscalar

|       | Integer instruction | FP instruction  | Clock cycle |
|-------|---------------------|-----------------|-------------|
| Loop: | LD F0,0(R1)         |                 | 1           |
|       | LD F6,-8(R1)        |                 | 2           |
|       | LD F10,-16(R1)      | ADDD F4,F0,F2   | 3           |
|       | LD F14,-24(R1)      | ADDD F8,F6,F2   | 4           |
|       | LD F18,-32(R1)      | ADDD F12,F10,F2 | 5           |
|       | SD 0(R1),F4         | ADDD F16,F14,F2 | 6           |
|       | SD -8(R1),F8        | ADDD F20,F18,F2 | 7           |
|       | SD -16(R1),F12      |                 | 8           |
|       | SD -24(R1),F16      |                 | 9           |
|       | SUBI R1,R1,#40      |                 | 10          |
|       | BNEZ R1,LOOP        |                 | 11          |
|       | SD -32(R1),F20      |                 | 12          |

- Unrolled 5 times to avoid delays (+1 due to SS)
- 12 clocks, or 2.4 clocks per iteration (1.5X)

9/22/99

CS252/Kubiatowicz  
Lec 8.23

## Dynamic Scheduling in Superscalar

- How to issue two instructions and keep in-order instruction issue for Tomasulo?
  - Assume 1 integer + 1 floating point
  - 1 Tomasulo control for integer, 1 for floating point
- Issue 2X Clock Rate, so that issue remains in order
- Only FP loads might cause dependency between integer and FP issue:
  - Replace load reservation station with a load queue; operands must be read in the order they are fetched
  - Load checks addresses in Store Queue to avoid RAW violation
  - Store checks addresses in Load Queue to avoid WAR, WAW
  - Called "decoupled architecture"

9/22/99

CS252/Kubiatowicz  
Lec 8.24

## Multiple Issue Challenges

- While Integer/FP split is simple for the HW, get CPI of 0.5 only for programs with:
  - Exactly 50% FP operations
  - No hazards
- If more instructions issue at same time, greater difficulty of decode and issue:
  - Even 2-scalar => examine 2 opcodes, 6 register specifiers, & decide if 1 or 2 instructions can issue
  - Multiported rename logic: must be able to rename same register multiple times in one cycle!
- VLIW: tradeoff instruction space for simple decoding
  - The long instruction word has room for many operations
  - By definition, all the operations the compiler puts in the long instruction word are independent => execute in parallel
  - E.g., 2 integer operations, 2 FP ops, 2 Memory refs, 1 branch
    - » 16 to 24 bits per field => 7\*16 or 112 bits to 7\*24 or 168 bits wide
  - Need compiling technique that schedules across several branches

9/22/99

CS252/Kubiatowicz  
Lec 8.25

## Loop Unrolling in VLIW

| Memory reference 1 | Memory reference 2 | FP operation 1  | FP op. 2        | Int. op/branch | Clock |
|--------------------|--------------------|-----------------|-----------------|----------------|-------|
| LD F0,0(R1)        | LD F6,-8(R1)       |                 |                 |                | 1     |
| LD F10,-16(R1)     | LD F14,-24(R1)     |                 |                 |                | 2     |
| LD F18,-32(R1)     | LD F22,-40(R1)     | ADDD F4,F0,F2   | ADDD F8,F6,F2   |                | 3     |
| LD F26,-48(R1)     |                    | ADDD F12,F10,F2 | ADDD F16,F14,F2 |                | 4     |
|                    |                    | ADDD F20,F18,F2 | ADDD F24,F22,F2 |                | 5     |
| SD 0(R1),F4        | SD -8(R1),F8       | ADDD F28,F26,F2 |                 |                | 6     |
| SD -16(R1),F12     | SD -24(R1),F16     |                 |                 |                | 7     |
| SD -32(R1),F20     | SD -40(R1),F24     |                 |                 | SUBI R1,R1,#48 | 8     |
| SD -0(R1),F28      |                    |                 |                 | BNEZ R1,LOOP   | 9     |

Unrolled 7 times to avoid delays

7 results in 9 clocks, or 1.3 clocks per iteration (1.8X)

Average: 2.5 ops per clock, 50% efficiency


Note: Need more registers in VLIW (15 vs. 6 in SS)

9/22/99

CS252/Kubiatowicz  
Lec 8.26

## Software Pipelining

- Observation: if iterations from loops are independent, then can get more ILP by taking instructions from different iterations
- Software pipelining: reorganizes loops so that each iteration is made from instructions chosen from different iterations of the original loop (- Tomasulo in SW)



9/22/99

CS252/Kubiatowicz  
Lec 8.27

## Software Pipelining Example

Before: Unrolled 3 times

```

1 LD F0,0(R1)
2 ADDD F4,F0,F2
3 SD 0(R1),F4
4 LD F6,-8(R1)
5 ADDD F8,F6,F2
6 SD -8(R1),F8
7 LD F10,-16(R1)
8 ADDD F12,F10,F2
9 SD -16(R1),F12
10 SUBI R1,R1,#24
11 BNEZ R1,LOOP
    
```


After: Software Pipelined

```

1 SD 0(R1),F4 ; Stores M[i]
2 ADDD F4,F0,F2 ; Adds to M[i-1]
3 LD F0,-16(R1); Loads M[i-2]
4 SUBI R1,R1,#8
5 BNEZ R1,LOOP
    
```

### Symbolic Loop Unrolling

- Maximize result-use distance
- Less code space than unrolling
- Fill & drain pipe only once per loop vs. once per each unrolled iteration in loop unrolling



9/22/99

CS252/Kubiatowicz  
Lec 8.28

## Software Pipelining with Loop Unrolling in VLIW

| Memory reference 1 | Memory reference 2 | FP operation 1  | FP op. 2 | Int. op/branch | Clock |
|--------------------|--------------------|-----------------|----------|----------------|-------|
| LD F0,-48(R1)      | ST 0(R1),F4        | ADDD F4,F0,F2   |          |                | 1     |
| LD F6,-56(R1)      | ST -8(R1),F8       | ADDD F8,F6,F2   |          | SUBI R1,R1,#24 | 2     |
| LD F10,-40(R1)     | ST 8(R1),F12       | ADDD F12,F10,F2 |          | BNEZ R1,LOOP   | 3     |

- **Software pipelined across 9 iterations of original loop**

– In each iteration of above loop, we:

- » Store to m,m-8,m-16 (iterations I-3,I-2,I-1)
- » Compute for m-24,m-32,m-40 (iterations I,I+1,I+2)
- » Load from m-48,m-56,m-64 (iterations I+3,I+4,I+5)

- **9 results in 9 cycles, or 1 clock per iteration**

- **Average: 3.3 ops per clock, 66% efficiency**

**Note: Need less registers for software pipelining (only using 7 registers here, was using 15)**

9/22/99

CS252/Kubiatowicz  
Lec 8.29

## Trace Scheduling

- **Parallelism across IF branches vs. LOOP branches**


- **Two steps:**

- **Trace Selection**

- » Find likely sequence of basic blocks (*trace*) of (statically predicted or profile predicted) long sequence of straight-line code

- **Trace Compaction**

- » Squeeze trace into few VLIW instructions
- » Need bookkeeping code in case prediction is wrong



- **Compiler undoes bad guess (discards values in registers)**

- **Subtle compiler bugs mean wrong answer vs. poorer performance; no hardware interlocks**

9/22/99

CS252/Kubiatowicz  
Lec 8.30

## Advantages of HW (Tomasulo) vs. SW (VLIW) Speculation

- HW determines address conflicts
- HW better branch prediction
- HW maintains precise exception model
- HW does not execute bookkeeping instructions
- Works across multiple implementations
- SW speculation is much easier for HW design

9/22/99

CS252/Kubiatowicz  
Lec 8.31

## Superscalar v. VLIW

- Smaller code size
- Binary compatibility across generations of hardware
- Simplified Hardware for decoding, issuing instructions
- No Interlock Hardware (compiler checks?)
- More registers, but simplified Hardware for Register Ports (multiple independent register files?)

9/22/99

CS252/Kubiatowicz  
Lec 8.32



## Intel/HP "Explicitly Parallel Instruction Computer (EPIC)"

- 3 Instructions in 128 bit "groups"; field determines if instructions dependent or independent
  - Smaller code size than old VLIW, larger than x86/RISC
  - Groups can be linked to show independence > 3 instr
- 64 integer registers + 64 floating point registers
  - Not separate files per functional unit as in old VLIW
- Hardware checks dependencies (interlocks => binary compatibility over time)
- Predicated execution (select 1 out of 64 1-bit flags) => 40% fewer mispredictions?
- IA-64 : instruction set architecture; EPIC is type
- Merced is name of first implementation (1999/2000?)
- LIW = EPIC?

9/22/99

CS252/Kubiatowicz  
Lec 8.33

## Limits to Multi-Issue Machines

- Inherent limitations of ILP
  - 1 branch in 5: How to keep a 5-way VLIW busy?
  - Latencies of units: many operations must be scheduled
  - Need about Pipeline Depth x No. Functional Units of independent operations to keep all pipelines busy.
  - Difficulties in building HW
  - Easy: More instruction bandwidth
  - Easy: Duplicate FUs to get parallel execution
  - Hard: Increase ports to Register File (bandwidth)
    - » VLIW example needs 7 read and 3 write for Int. Reg. & 5 read and 3 write for FP reg
  - Harder: Increase ports to memory (bandwidth)
  - Decoding Superscalar and impact on clock rate, pipeline depth?

9/22/99

CS252/Kubiatowicz  
Lec 8.34

## Limits to Multi-Issue Machines

- Limitations specific to either Superscalar or VLIW implementation
  - Decode issue in Superscalar: how wide practical?
  - VLIW code size: unroll loops + wasted fields in VLIW
    - » IA-64 compresses dependent instructions, but still larger
  - VLIW lock step => 1 hazard & all instructions stall
    - » IA-64 not lock step? Dynamic pipeline?
  - VLIW & binary compatibility IA-64 promises binary compatibility

9/22/99

CS252/Kubiatowicz  
Lec 8.35

## Limits to ILP

- Conflicting studies of amount
  - Benchmarks (vectorized Fortran FP vs. integer C programs)
  - Hardware sophistication
  - Compiler sophistication
- How much ILP is available using existing mechanisms with increasing HW budgets?
- Do we need to invent new HW/SW mechanisms to keep on processor performance curve?
  - Intel MMX
  - Motorola AltaVec
  - Supersparc Multimedia ops, etc.

9/22/99

CS252/Kubiatowicz  
Lec 8.36

## Limits to ILP

Initial HW Model here; MIPS compilers.

Assumptions for ideal/perfect machine to start:

1. **Register renaming**-infinite virtual registers and all WAW & WAR hazards are avoided
2. **Branch prediction**-perfect; no mispredictions
3. **Jump prediction**-all jumps perfectly predicted => machine with perfect speculation & an unbounded buffer of instructions available
4. **Memory-address alias analysis**-addresses are known & a store can be moved before a load provided addresses not equal


1 cycle latency for all instructions; unlimited number of instructions issued per clock cycle

9/22/99

CS252/Kubiatowicz  
Lec 8.37

## Upper Limit to ILP: Ideal Machine

(Figure 4.38, page 319)




9/22/99

CS252/Kubiatowicz  
Lec 8.38

## More Realistic HW: Branch Impact

Figure 4.40, Page 323




9/22/99

Perfect Pick Cor. or BHT BHT (512) Profile No prediction

CS252/Kubiatowicz  
Lec 8.39

## More Realistic HW: Register Impact

Figure 4.44, Page 328




9/22/99

Infinite 256 128 64 32 None

CS252/Kubiatowicz  
Lec 8.40

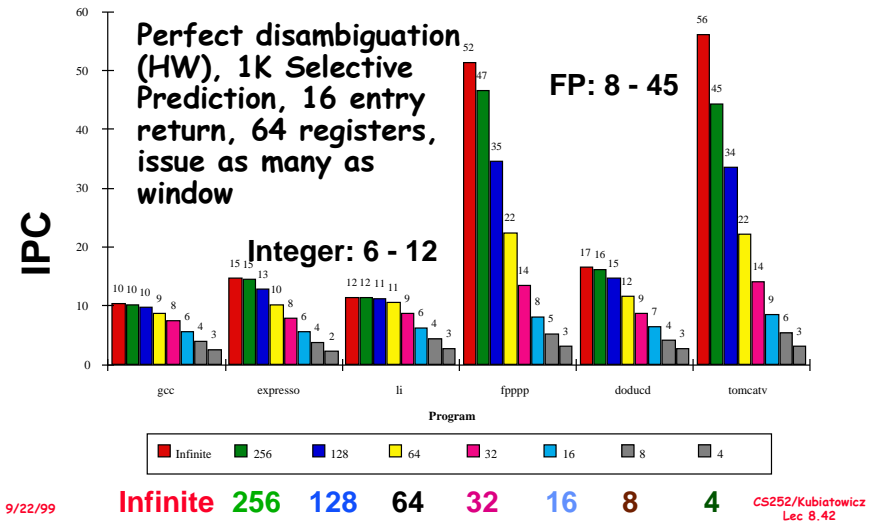
## More Realistic HW: Alias Impact

Figure 4.46, Page 330




## Realistic HW for '9X: Window Impact

(Figure 4.48, Page 332)



## Braniac vs. Speed Demon(1993)

- 8-scalar IBM Power-2 @ 71.5 MHz (5 stage pipe)
- vs. 2-scalar Alpha @ 200 MHz (7 stage pipe)




## Problems with scalar approach to ILP extraction

- Limits to conventional exploitation of ILP:
  - pipelined clock rate: at some point, each increase in clock rate has corresponding CPI increase (branches, other hazards)
  - instruction fetch and decode: at some point, its hard to fetch and decode more instructions per clock cycle
  - cache hit rate: some long-running (scientific) programs have very large data sets accessed with poor locality; others have continuous data streams (multimedia) and hence poor locality

## Preview: Alternative Model: Vector Processing

- Vector processors have high-level operations that work on linear arrays of numbers: "vectors"



9/22/99

CS252/Kubiatowicz  
Lec 8.45

## Properties of Vector Processors

- Each result independent of previous result  
=> long pipeline, compiler ensures no dependencies  
=> high clock rate
- Vector instructions access memory with known pattern  
=> highly interleaved memory  
=> amortize memory latency of over - 64 elements  
=> no (data) caches required! (Do use instruction cache)
- Reduces branches and branch problems in pipelines
- Single vector instruction implies lots of work (- loop)  
=> fewer instruction fetches

9/22/99

CS252/Kubiatowicz  
Lec 8.46

## Vector Advantages

- Easy to get high performance; N operations:
  - are independent
  - use same functional unit
  - access disjoint registers
  - access registers in same order as previous instructions
  - access contiguous memory words or known pattern
  - can exploit large memory bandwidth
  - hide memory latency (and any other latency)
- Scalable (get higher performance as more HW resources available)
- Compact: Describe N operations with 1 short instruction (v. VLIW)
- Predictable (real-time) performance vs. statistical performance (cache)
- Multimedia ready: choose N \* 64b, 2N \* 32b, 4N \* 16b, 8N \* 8b
- Mature, developed compiler technology
- Vector Disadvantage: Out of Fashion?

9/22/99

CS252/Kubiatowicz  
Lec 8.47

## Summary

- Dynamic hardware schemes can unroll loops dynamically in hardware
- Explicit Renaming: more physical registers than needed by ISA. Uses a translation table
- Precise exceptions/Speculation: Out-of-order execution, In-order commit (reorder buffer)
- Superscalar and VLIW: CPI < 1 (IPC > 1)
  - Dynamic issue vs. Static issue
  - More instructions issue at same time => larger hazard penalty
  - Limitation is often number of instructions that you can successfully fetch and decode per cycle => "Flynn barrier"
- SW Pipelining
  - Symbolic Loop Unrolling to get most from pipeline with little code expansion, little overhead

9/22/99

CS252/Kubiatowicz  
Lec 8.48