

CSc 553 — Principles of Compilation

24 : Instruction Scheduling I

Christian Collberg
Department of Computer Science
University of Arizona
collberg@gmail.com

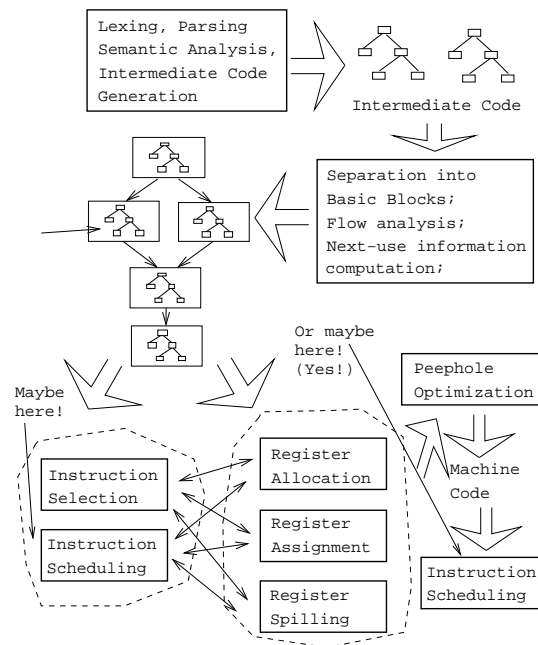
Copyright © 2011 Christian Collberg

March 24, 2011

1

Introduction

2



3 Instruction Scheduling

- Instruction scheduling is an inherently **machine dependent** task. Unlike other code improvements which work with high- or medium-level intermediate representations, instruction scheduling works explicitly with machine instructions.

- Instruction scheduling is often one of the last phases of code generation, performed after instruction selection and register allocation.
- However, scheduling affects register allocation, so, ideally, we'd like to do instruction scheduling and register allocation simultaneously.
- We'll look at a scheduling algorithm that works directly on the generated machine code.

4

Instruction Pipelines

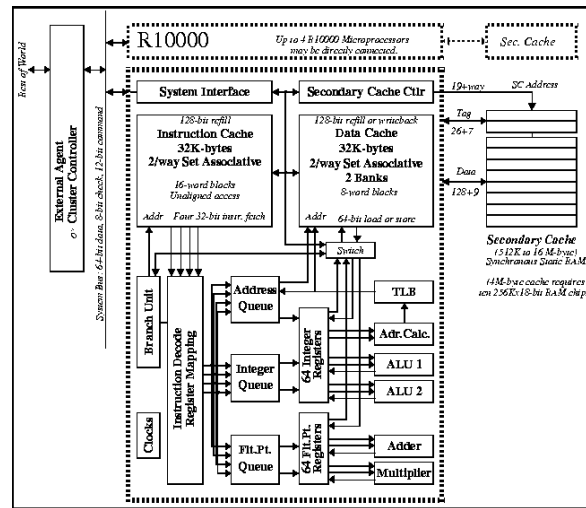
5 Instruction Pipelines

- Modern processors use instruction pipelines to speed up execution **throughput** (total amount of work done in a given time).
- The idea is that the execution of several instructions are overlapped in time, much the same as when a number of cars are assembled on an assembly line.
- The execution of instructions is broken down in several (4–7) individual pipeline stages.
- Pipelining doesn't speed up the execution of individual instructions; rather, it increases the number of simultaneously executing instructions and the rate at which they are started and completed.

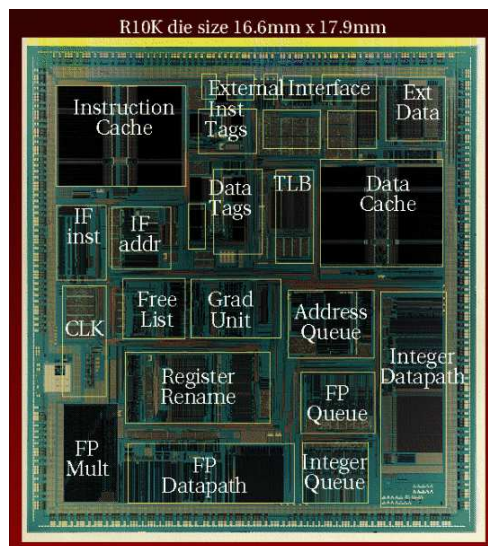
6 Instruction Pipelines II

- A typical 5-stage pipeline:
 - (1) **Fetch** Fetch the instruction from memory (or instruction cache).
 - (2) **Decode** Decode the instruction and fetch the values of registers that the instruction will need.
 - (3) **Execute** Execute the instruction.
 - (4) **Memory** Access memory (for load and store instructions).
 - (5) **Write** Write the result of the instruction into the result register.
- Not all instructions make use of all the pipeline stages. There is often more than one execute stage.
- The MIPS R10000 has 7 stages: (1) Fetch, (2) Decode, (3) Issue Instruction and Read Registers, (4) Exec 1, (5) Exec 2, (6) Exec 3, (7) Register Write.

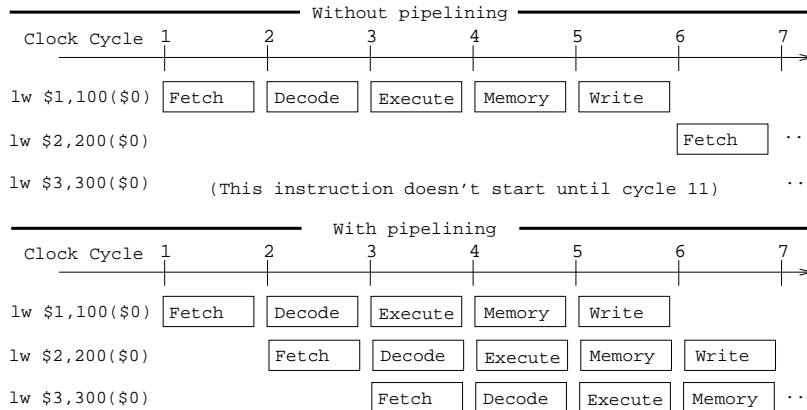
7 MIPS R10000 Block Diagram



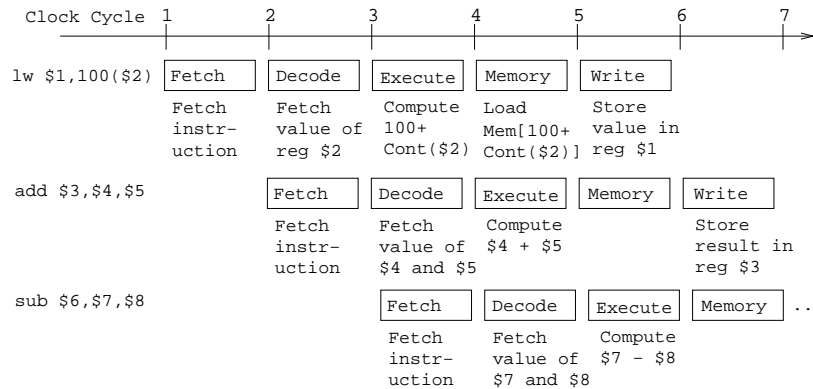
8 MIPS R10000 Die Photo



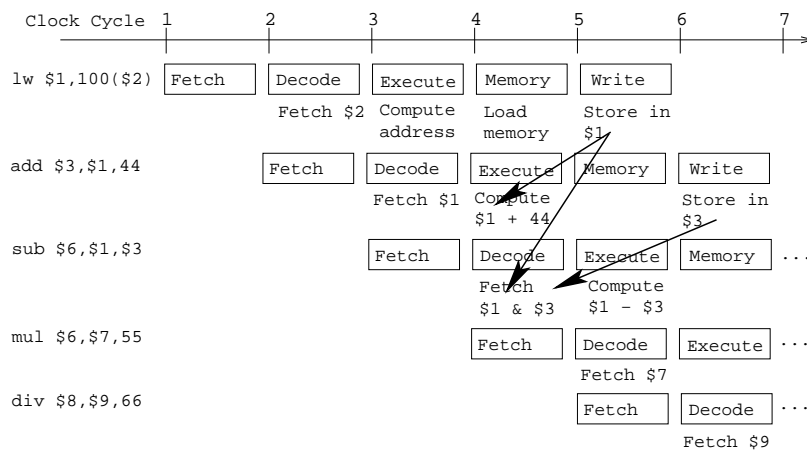
9 Pipeline Example I



10 Pipeline Example II



11 Pipeline Data Hazards I



12 Pipeline Data Hazards II

- The previous example shows one kind of problem we have with pipelines. Data hazards occur when the result of one instruction isn't ready in time for when that result is needed.

- Some processors will detect such situations and **stall** the pipeline until the value is ready. This is called a **hardware/pipeline interlock**.
- Interlocks are expensive (extra control logic on the chip which takes up valuable chip real estate) so some processors do away with them. Instead they rely on compilers and assemblers to insert NOPs when needed.

13 Pipeline Data Hazards III

- Example: On a MIPS 2000, the value loaded by a `ld` instruction isn't available to the immediately following instruction. The processor doesn't have hardware interlocks. If you give the following code to the assembler

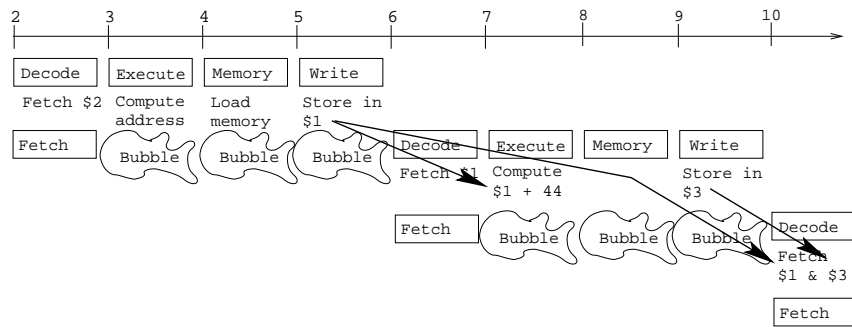
```
ld      100($4), $5
add     $5, $5, 56
```

it will insert the necessary NOPs:

```
ld      100($4), $5
nop
add     $5, $5, 56
```

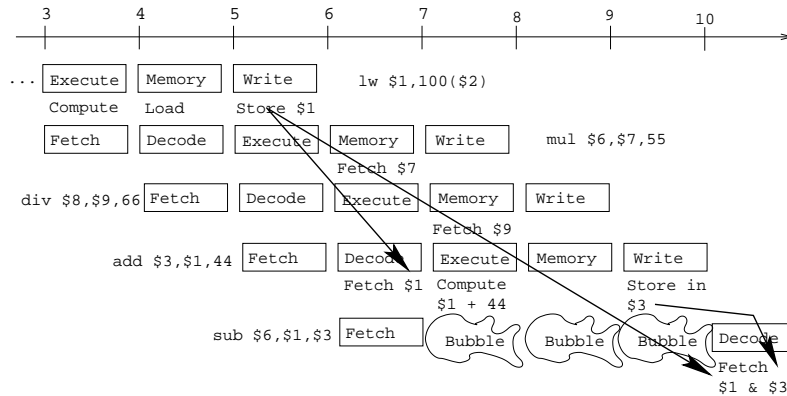
- Even if the hardware does have interlocks, the compiler (or assembler) should pay attention to the order in which instructions are scheduled. A well scheduled program may be up to 50% faster than a naively scheduled one.

14 Pipeline Data Hazards – Stalls



- Some processors take care of data hazards themselves – the pipeline is **stalled** for a number of cycles until the hazard is resolved. This is like inserting one or more **bubbles** (NOPs) in the pipeline.

15 Pipeline Data Hazards – Reordering



- The compiler can sometimes reorder instructions to avoid stalls.

16

Dependency Graphs

17 Dependency Graphs I

- An instruction scheduler reorders the instructions in an attempt at minimizing pipeline stalls.
- Obviously, we must not rearrange the code so that (when executed) it produces a different result from before!

```
add $3,$2,44  Wrong!  li $2,66
li $2,66      =>      add $3,$2,44
```

- We must therefore model all the dependencies between the instructions. We store this information in a directed **dependency graph**. The nodes of the graph are the instructions we're scheduling, and there's an edge $a \Rightarrow b$ if instruction a must come before instruction b .

18 Dependency Graphs II

- There can be three kinds of dependencies between instructions:

flow dependence

- Also, **true dependence** or **definition-use dependence**.

(i) X := ...

(j) ... := X

- Instruction (i) generates (**defines**) a value which is used by instruction (j). We write (i) \longrightarrow (j).

19 Dependency Graphs III

anti-dependence

- Also, **use-definition dependence**.

(i) $\dots := X$

(j) $X := \dots$

- Instruction (i) uses a value overwritten by instruction (j). We write $(i) \rightarrow (j)$.

20 Dependency Graphs IV

Output-dependence

- Also, **definition-definition dependence**.

(i) $X := \dots$ \dots

(j) $X := \dots$

- Instructions (i) and (j) both assign to (**define**) the same variable. We write $(i) \ominus (j)$.
- Regardless of the type of dependence, if instruction (j) depends on (i), then (i) has to be scheduled before (j).

21 Dependency Graphs V

Flow Dependence Example

(1) `ld 100($3), $2`
 $\dots \Rightarrow (1) \rightarrow (5)$

(5) `add $5, $2, 55`

- We can't move (1) after (5). If we did, \$2 would not have the correct value at (5). Hence there is a flow dependence between (1) and (5).

Anti-Dependence Example

(2) `add $5, $2, 55`
 $\dots \Rightarrow (2) \rightarrow (6)$

(6) `li $2, 33`

- \$2 gets a new value in (6). If we moved (2) after (6) \$2 would have the wrong value in (6).

22 Dependency Graphs VI

Output-Dependence Example

(3) `li $2, 33`

(4) `add $2, $5, 55` $\Rightarrow (3) \ominus (4)$

(5) `add $6, $2, 44`

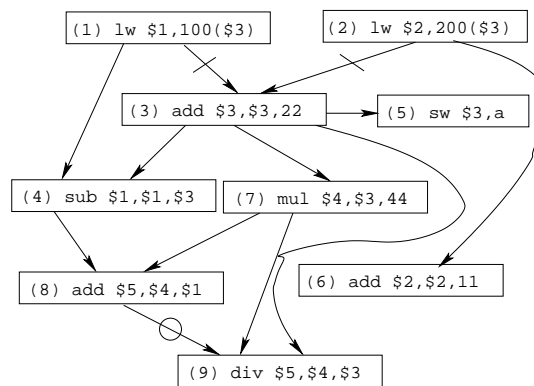
- If we switched (3) and (4), \$2 would get the wrong value in (3).

23 Dependency Graphs VII

Complex Example

```
(1) lw    $1, 100($3)
(2) lw    $2, 200($3)
(3) add   $3, $3, 22
(4) sub   $1, $1, $3
(5) sw    $3, 500($0)
(6) add   $2, $2, 11
(7) mul   $4, $3, 44
(8) mul   $5, $4, $1
(9) div   $5, $4, $3
```

24 Dependency Graphs VIII



- If instruction (b) depends on instruction (a) then (b) can be scheduled after a delay of one cycle.
- Real pipelines are more complex than this....

25

Topological Sorting

26 Topological Sorting I

- Problem: How to get dressed? I can't put my clothes on in an arbitrary order: socks have to be put on before shoes, shirt before tie, and so on.
- Each garment becomes a node in a DAG, and there's an edge from u to v if I have to put on u before v .

Topological Sorting:

“Order the nodes of the graph G in a sequence such that if $x \rightarrow y$ is an edge in G then x comes before y in the ordering.”

Simple Algorithm:

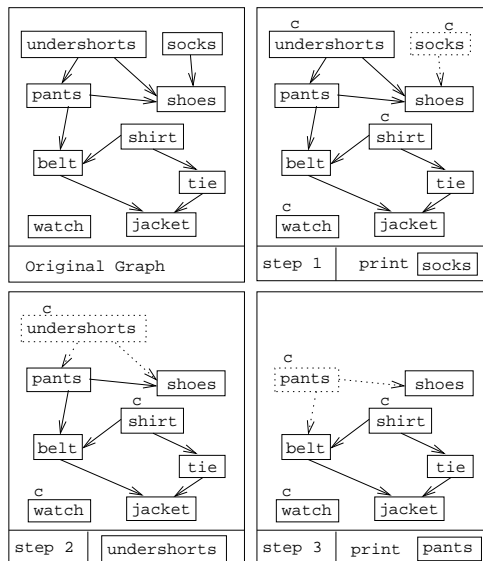
Repeat until no more nodes:

- Pick a node n without predecessors.
- Print n .
- Delete n and all its outgoing edges.

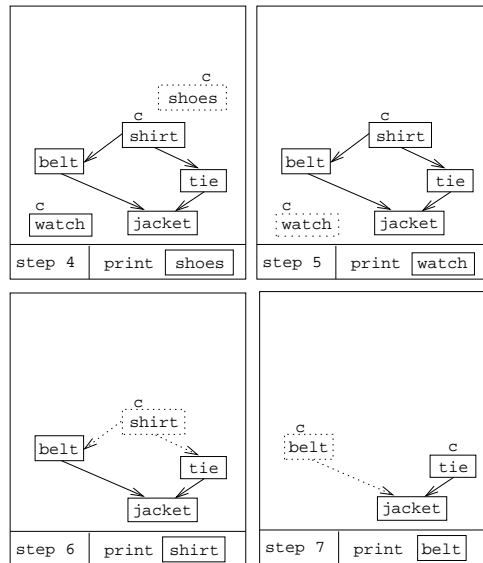
27 Topological Sorting II

- In the following example, candidates are marked with a c .
- When there is more than one candidate to choose from, we pick one at random.
- There are often many possible topological orders to choose from. For example, since the watch node has no dependencies at all, I can put it on at any time.

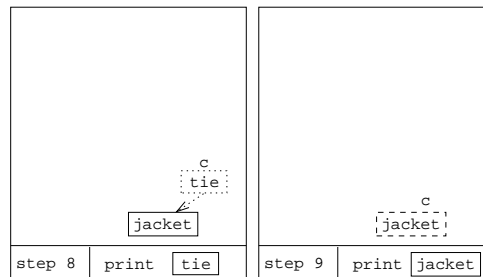
28 Topological Sorting III



29 Topological Sorting IV




30 Topological Sorting V



31

Complete Topological Order



Algorithm

32

```
Last := {}; /* Last scheduled instr. */
REPEAT
```

```
  Candidates := set of all nodes without
    predecessors (incoming edges);
```

```
  Colliding := set of all instructions that collide with
    the instructions in Last; i.e. nodes  $n$  for which
    there is an arch Last  $\rightarrow n$ .
```

```

realCandidates:= Candidates-Colliding;
IF realCandidates ≠ {} THEN
  b := Use heuristic to pick the best
      realCandidate to schedule;
  Remove b and all its outgoing edges from the graph;
  print b; Last := {b};
ELSE
  print NOP; Last := {};
ENDIF
UNTIL Candidates = {};

```

33 Heuristic

- In the “standard” topological sorting algorithm we pick a node at random when there are several candidate nodes to choose from. In our scheduling algorithm we need a better heuristic.
- If there is more than one real candidate (a candidate that doesn’t collide with the last instruction scheduled) we choose one according to these criteria:
 1. Pick the candidate with the largest number of outgoing edges. Scheduling this instruction early will hopefully give us more freedom to schedule the remaining instructions.
 2. If all candidates have the same number of outgoing edges, pick the one that has the longest path to a leaf.

34

Example

35 Example I (a)

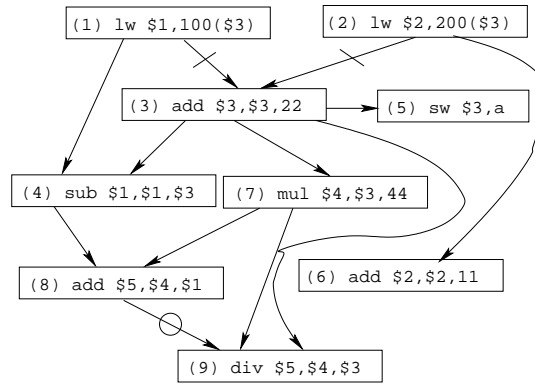
Complex Example

```

(1)  lw    $1, 100($3)
(2)  lw    $2, 200($3)
(3)  add   $3, $3, 22
(4)  sub   $1, $1, $3
(5)  sw    $3, 500($0)
(6)  add   $2, $2, 11
(7)  mul   $4, $3, 44
(8)  mul   $5, $4, $1
(9)  div   $5, $4, $3

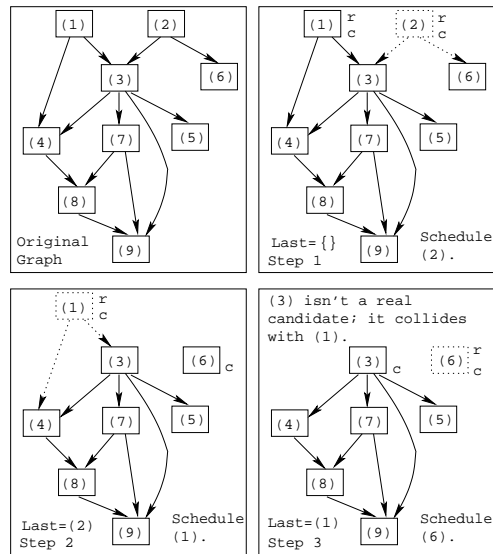
```

36 Example I (b)

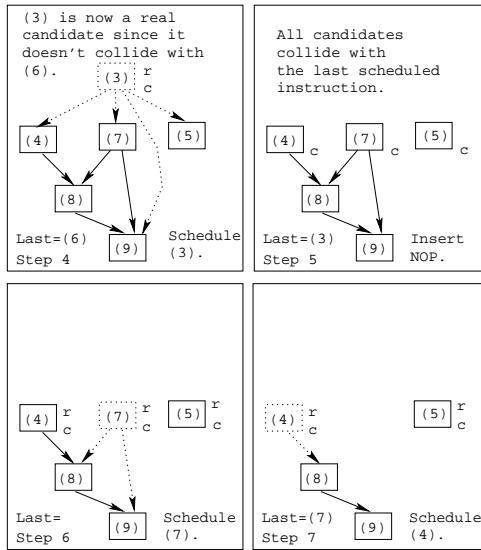


- If instruction (b) depends on instruction (a) then (b) can be scheduled after a delay of one cycle.
- We mark real candidates with an r.

37 Example I(c)



38 Example I(d)



39 Example I(e)

Last=(4) Step 8 Schedule (5).

Last=(8) Step 9 Schedule (8).

Last=(8) Step 10 Schedule (9).

Generated Code

```

(2) lw $2,200($3)
(1) lw $1,100($3)
(6) add $2,$2,11
(3) add $3,$3,22
NOP
(7) mul $4,$3,44
(4) sub $1,$1,$3
(5) sw $3,a
(8) add $5,$4,$1
(9) div $5,$4,$3
                
```

NOTE:
(8) → (9)
is an output
dependence. Hence
there's no need for
a NOP between them.

40

Summary

41 Readings and References

- Read the Tiger book, Chapter 20, pp. 474–497.
- For a background on pipelining and hazards, see Patterson & Hennessy, “Computer Organization and Design – The Hardware/Software Interface”, pp. 364–367,

- Wilhelm & Maurer, “Compiler Design”, pp. 557–558, 570–580.
- Steven Muchnick, *Advanced Compiler Design and Implementation*, Section 9.2, pp. 269–274 and Chapter 17, pp. 531–547.

42 Summary

- Some processors have instructions with pipeline **hot-spots**: A value that is computed in one cycle must be consumed in the next cycle (before it is overwritten).
- We have assumed that the delay between two instructions is always exactly 1 (one) cycle. Real processors require different delays depending on the instruction.

43

Homework

44 Homework I

- Consider the following basic block:

```
(1) lw    $1, 100($2)
(2) li    $3, 100
(3) add   $2, $3, 55
(4) sub   $1, $3, 66
(5) li    $5, 99
(6) add   $7, $3, $5
(7) li    $3, 93
(8) add   $9, $2, $7
(9) sub   $1, $9, $3
```

1. Construct the dependency graph.
2. Apply the scheduling algorithm to the graph. Assume that there is a one-cycle delay between an instruction and when it’s computed value can be used.
3. Is the generated schedule an optimal one?

45

Exam Problems

46 Exam Problem I (415.730 ’96)

1. What is a **pipeline data hazard**?
2. Why is it important for the compiler to perform aggressive instruction scheduling, even when the architecture has **hardware interlocks**?
3. Describe the algorithms and data structures needed to perform instruction scheduling in the presence of pipeline data hazards.