1.1 Introduction

The material on dynamic scheduling is not covered in detail in the text, which is unfortunate since as of this writing most processors are statically scheduled. This study guide provides a summary of dynamic scheduling and a guide to sample problems from old homeworks and exams. (The solutions are sometimes detailed.)

In a statically scheduled processor instructions start execution in program order while in a dynamically scheduled processor instructions can start execution out of order. Statically scheduled systems are much simpler since instructions march through the pipeline in step, with bubbles (gaps) inserted where necessary. A shortcoming is that an instruction that must wait blocks all of the instructions behind it. Those instructions with dependencies on the waiting instruction would have to wait anyway but there is no reason to block other instructions, other than hardware cost. In a dynamically scheduled system the only instructions that normally must wait are those for which input operands are not ready. (There are other reasons for waiting, for example, the needed functional unit is busy.) Dynamically scheduled processors are far more costly but achieve better performance than static scheduling on superscalar processors, especially when loads can take more than a cycle or two (due to a cache miss, covered later).

As of this writing most processors are dynamically scheduled. This includes the Pentium III and Pentium 4, Alpha 21264, MIPS R10000, PowerPC 620, and HP-PA 8000, and the later versions of these processors. Two exceptions are the Sun UltraSparc III and the Intel Itanium 2, which are statically scheduled.

1.2 Summary of Dynamic Scheduling Method 3

In class three methods of dynamic scheduling were mentioned, but (in Spring 2003) only one was covered in detail, Method 3. In Method 3 values are stored in a physical register file and physical register numbers are used to re-name registers. This is the only method covered in the Spring 2003 semester and so Spring 2003 final exam questions are unlikely to use the other methods. The following is a brief description of Method 3. First, activities in each stage are covered, then the tables and other elements are described. The descriptions below refer to the following illustration:



1.2.1 Stage: IF, Instruction Fetch

This occurs strictly in program order. The maximum number of instructions that can be fetched per cycle is equal to the decode width. (An x-way superscalar processor has a decode width of x.)

For simplicity IF is shown using the same hardware as the one-way statically scheduled processor. In real dynamically scheduled superscalar systems additional hardware is needed for branch and target prediction and for shifting and masking data retrieved from the memory port (really a cache port).

1.2.2 Stage: ID, Instruction Decode

This occurs strictly in program order. The maximum number of instructions that can be decoded per cycle is equal to the decode width. The following is done:

An entry for the instruction is placed in the reorder buffer with the following information: The address of the instruction (PC); two status bits C, complete, and X, raised an exception; the architected destination register number, dst, (*e.g.*, t1, s2); the physical destination register, dstPR, (*e.g.*, p90,p103); and the incumbent physical register number, incumb.

The *incumbent* is the physical register holding the old value of the destination register. For example, consider the sub instruction in the example below. The physical registers used for the destinations are shown in the comments. Physical register 92 is removed from the free list and will be used for t1, the architected register holding the result of the subtract. The "old" value of t1 is in physical register 90, so 90 is the incumbent.

```
add t1, t2, t3 # t1 to be stored in physical register 90
or t5, t1, t7 # t5 to be stored in physical register 91
sub t1, t8, t5 # t1 to be stored in physical register 92
```

Also in DI, the architected source register numbers are translated into physical registers using the ID register map. The source physical register numbers are labeled rsPR and rtPR in the diagram.

A physical register is removed from the free list and assigned to the destination. This new physical register is labeled dstPR on the diagram.

The new physical register is written into the ID map (where following instructions will find it) and the incumbent, incmb, is retrieved. (The incumbent is overwritten by the new physical register number.) As described above the incumbent is written into the reorder buffer.

Control determines which instruction queue to put the instruction in (IQ in the diagram) and which operation to perform (Op). (The diagram only shows one instruction queue.) The reorder buffer provides a unique number associated with the new ROB entry, ROB #, this is used during WB.

1.2.3 Stage: Q, Instruction Queue

The Q stage can refer to two things: being placed in the instruction queue, QI, or being removed from it QO. It is used both ways in the solutions (just Q) and so both are correct in answers. QI occurs strictly in program order, QO does not necessarily occur in program order.

An instruction is placed in the instruction queue, QI, in the cycle after ID (unless there is a stall, which is very rare in the homework and exam problems).

An instruction is removed from the queue, QO, when it is ready to execute (its operands will be available in the next cycle). When it is removed from the queue it reads physical register values (rsVal, rtVal) from the physical register file.

After Q the instruction moves to a *functional unit*. The following functional units are common in the homework and exams: EX, integer and logic operations; B, branch resolution; L1 L2 load/store unit (discussed further below), A1 A2... floating-point add, M1 M2... floating-point multiply.

1.2.4 Stages: EX, A1, M1, etc. Arithmetic and Logical Functional Units

These work like their statically scheduled counterparts. By default a complete set of bypass connections is assumed. After exiting a functional unit the instructions go to writeback, WB.

1.2.5 Stage: B, Branch Resolution Functional Unit

This is used to determine if a branch is taken or not (the *outcome*). See WB for more on recovery.

1.2.6 Stages: L1, L2, Load/Store Unit

In the L1 step the effective address of the load or store is computed. For loads L2 indicates that the load/store queue and, if necessary, the cache is being checked for the data. If the data is found the next segment is WB, otherwise no segment is shown until the data arrives at which time L2 is shown again. Stores first write their data into the load/store queue (L2) where following loads can read it. When they commit the data is written to the cache.

In the example below the lw hits the cache and lh misses the cache. lw \$t1, 0(\$t2) IF ID Q L1 L2 WB # Cache hit add \$t3, \$t1, \$t4 IF ID Q Q EX WB lh \$t1, 0(\$t2) IF ID Q L1 L2 L2 WB # Cache miss. add \$t3, \$t1, \$t4 IF ID Q Q EX WB

1.2.7 Stage: WB, Writeback

The WB stage occurs after the last functional unit stage (EX, L2, M6, etc). The default assumption for all problems is that any number of write-backs can be performed per cycle. This assumption is unrealistic but it makes solving problems far less tedious.

The following occurs during writeback: The result is written to the physical register file. The status of the instruction (complete, and whether an exception occurred) is written to the reorder buffer.

If the instruction is a branch and if the ID map has been backed up, the outcome of the branch (taken, not taken) is compared to the predicted outcome. If they differ recovery starts. Recovery consists of copying the commit map to the ID map and restoring the free list. If the ID map has not been backed up recovery will wait until the branch commits.

1.2.8 Stage: C, Commit

Commit occurs strictly in program order. By default the maximum number of instructions that can commit per cycle is equal to the decode width.

During commit the commit register map is updated and the incumbent is put back on the free list. Remember, the incumbent is **not** the physical register assigned to the committing instruction, it is the physical register assigned to the last instruction that wrote the same architected register. (See the example for the ID stage.)

1.2.9 Reorder Buffer

A queue holding *in-flight* instructions. Instructions always enter the ROB in ID and normally leave in C. They can also be *flushed*, squashed *en-masse*, if there is a recovery due to an exception or misprediction.

An instruction will update its ROB entry during WB, indicating that it has completed and whether it has raised an exception.

If over a long enough period instructions are fetched faster than they are being committed the ROB will fill up, stalling IF and ID. This can affect the CPI in some problems, such as Spring 2003 Homework 6, Problem 1.

1.2.10 Instruction Queue/Scheduler

A list of instructions that have not yet executed. (It's called a queue but it's not first-in/first-out.) The scheduler monitors which functional units are busy and which registers are ready. From this information it chooses instructions to start, they will enter the QO (usually just shown as Q) stage.

Real systems may use several schedulers, each scheduler can send instructions to a subset of functional units. For example, there may be an integer instruction queue and a floating-point instruction queue.

1.2.11 ID Map

A table giving the latest physical register assigned to each architected register. Written in ID; the entry number that is written is based on the architected register, that entry is written with the new physical register number. Since it's written in the ID stage it reflects the part of the program that has passed through ID.

1.2.12 Commit Map

A table giving the latest physical register assigned to each architected register. Written in C; the entry number that is written is based on the architected register, that entry is written with the new physical register number. Since it's written in the commit stage it reflects the part of the program that has passed through commit.

1.2.13 Physical Register File

A table giving register values. It is written in the WB stage using, of course, physical register numbers.

1.2.14 Free List

A list of unused physical registers. An instruction removes a physical register when it passes through ID. That physical register will hold the result of the instruction. If an instruction were proud of its creation (the result) then it would hate to see it tossed into the recycle bin (free list). It is very fortunate that instructions do not through their creations into the free list, but instead throw out their predecessor's creation. (That is, they put the incumbent in the free list.)

1.3 Problems, By Type

Problem types are described below in roughly order of increasing difficulty.

1.3.1 Pipeline Execution Diagram

Show a pipeline execution diagram for the code fragment below running on the following system ... Solving this type of problem is a matter of knowing the steps for instruction execution.

Solution Tips: Be sure to check for dependencies. Remember that instructions execute as soon as their source operands are ready (unless there is not a free functional unit, which is rare for this type of problem). If there is a branch misprediction, determine whether the ID Map is backed up; if it is recovery starts in WB, otherwise it starts during commit.

See the following final exam problems: Spring 2000 problem 3, Spring 2001 problem 2, 1998 problem 2, 1997 problem 2.

1.3.2 Complete ID Map and Other Tables

A program executes on a dynamically scheduled system using method 3 as shown in the pipeline execution diagram below. Using the tables provided, show the changes to the ID Map, Commit Map, and Physical Register File. Also show when each instruction commits.

See the following final exam problems: Spring 2002 problem 1, Fall 2001 problem 1, Spring 2000 problem 2.

1.3.3 Load/Store Unit

See Fall 2001 problem 4a.

1.3.4 Non-Standard Systems

Many of the problems above could be solved by memorizing rote procedures. To test true understanding some problems ask about an unusual (not covered in class) dynamically scheduled system. There are two examples of this sort of problem (so far). In the first (Fall 2003 problem 2) the hardware is defective; tables must be filled in showing the incorrect execution. In the second, Spring 2002 problem 2 (the later parts) asks about the execution of predicated instructions on a dynamically scheduled system.