

Diplomarbeit

Design and Evaluation of a Superscalar RISC Processor

Mark A. Hillebrand
(mah@wjpserver.cs.uni-sb.de)

Lehrstuhl Prof. Dr. W. J. Paul
Fachbereich Informatik
Universität des Saarlandes

April 2000

Eidesstattliche Erklärung

Hiermit erkläre ich an Eides Statt, daß ich die vorliegende Arbeit selbständig und ohne unerlaubte fremde Hilfe angefertigt und andere als die angegebenen Quellen und Hilfsmittel nicht benutzt habe.

Saarbrücken, den 5. April 2000

Contents

1	Introduction	9
1.1	Overview	9
1.2	Notes on Reading	10
2	Abstract Model	11
2.1	Tomasulo Algorithm	11
2.1.1	Sequential Instruction-Stream-Based Machine	12
2.1.2	Tomasulo Machine	12
2.1.3	Correctness	19
2.2	Roll-Back	22
2.2.1	Definition of Precise Roll-Back	23
2.2.2	Roll-Back Retirement Protocol	23
2.2.3	Proof of Preciseness	24
2.3	Instruction Fetch and Speculation	24
2.3.1	Sequential Instruction-Memory-Based Machine	24
2.3.2	Non-Speculative Instruction Fetch	25
2.3.3	Speculative Instruction Fetch	27
3	Hardware	31
3.1	Notation	31
3.1.1	Control Signals	31
3.1.2	Busses	33
3.1.3	Figures	33
3.2	Half-Unary Number Format	34
3.3	Instruction Fetch Mechanism	36
3.3.1	Overview	36
3.3.2	Algorithms	37
3.3.3	Control	40
3.4	Instruction Memory Environment	44
3.5	Instruction Fetch Queue Environment	44
3.6	PC environment	45
3.7	Prediction Environment	47
3.7.1	Finding the first CFI	47
3.7.2	Prediction	48
3.7.3	Resolving	48
3.7.4	Construction of the CFI bus	49
3.8	Instruction Window Environment	49
3.8.1	Construction of the Instruction Window	49
3.8.2	Draining and Switching the Issuing IFQ	50
3.9	Decode / Issue Environment	51
3.9.1	Decoding the Instruction	51
3.9.2	Issuing the Instruction	52

3.10	Reservation Station Environment	60
3.10.1	Scheduling of the reservation station input busses	60
3.10.2	Snooping for source operands	61
3.11	Function Unit Environments	61
3.11.1	Arithmetical and Logical Functional Unit	61
3.11.2	Floating Point Functional Units	62
3.11.3	Data Memory Functional Unit	62
3.11.4	Branch Checker Unit	65
3.11.5	Producer Environments	67
3.12	CDB Control Environment	68
3.13	Reorder Buffer Environment	68
3.13.1	ROB Queue Control	70
3.13.2	Issue	70
3.13.3	Forwarding	70
3.13.4	Completion	71
3.13.5	Retirement	71
3.14	Register File Environment	73
3.14.1	General Purpose Register File	74
3.14.2	Floating Point Register File	74
3.14.3	Special Purpose Register File	75
3.15	Producer Table Environment	78
3.15.1	General Purpose Register Producer Table	78
3.15.2	Floating Point Register Producer Table	80
3.15.3	Special Purpose Register Producer Table	81
4	Evaluation	85
4.1	Hardware Model	85
4.2	Parameter Space	85
4.3	A 2-Superscalar Processor	86
4.3.1	Cost and Delay Optimization	88
4.3.2	Longest Path	91
4.3.3	Quality Comparison	91
5	Circuits	93
5.1	Multiported Round-Robin Selector	93
5.1.1	Abstract View on Multiported Round-Robin Selectors	93
5.1.2	Implementation of a Multiported Round-Robin Selector	94
5.2	Multicounter	96
5.3	Multiported Queue	96
5.3.1	Definition of an Abstract Multiported Queue	96
5.3.2	Interface	98
5.3.3	Register-Based Implementation	99
5.3.4	RAM-Based Implementation	103
5.4	Reservation Station Queue	105
5.4.1	Definition of an Abstract Multiported Queue	105
5.4.2	Description of the Interface and Equivalence Criterion	106
5.4.3	Register-Based Implementation	106
6	Perspective	111
A	DLX Instruction Set Architecture	113

B	Sample Branch Predictor Unit	117
B.1	Branch Predictor	117
B.1.1	Implementation	117
B.1.2	Integration in the processor	119
B.2	Write-Buffered RAM	119
B.2.1	Write	119
B.2.2	Read	121
C	Auxiliary Circuits	123
C.1	Find-First-One Half-Unary	123
C.1.1	Definition and Construction	123
C.1.2	Correctness.	123
C.2	Find-First- k -Ones	123
C.2.1	Definition	123
C.2.2	Construction	124
C.2.3	Correctness	124
C.2.4	Different Interpretation	124
C.3	Multiple Incrementer	125

Chapter 1

Introduction

Processor design tries to increase the computing power of microprocessors by the advances in two fields of research. In the first, speed-up is achieved by physical advances in circuitry like reduced gate delay and an increased integration density or wafer size. In the second, refined or new algorithmic strategies are implemented decreasing the cycle-per-instruction (CPI) rate of a processor.

One approach in the second field that has been followed with considerable success is the development of superscalar processors. Such processors execute instructions in parallel while retaining the traditional sequential semantics of programs.

This thesis develops a superscalar processor on formal basis. Our superscalar processor falls into two parts: an instruction fetch mechanism and a superscalar DLX processor core implementing the Tomasulo algorithm. The instruction fetch mechanism loads instruction from the instruction memory and by sorting them in program order produces the so-called *instruction stream*. The processor takes the instruction stream and executes it as fast as possible. This combination of fetch mechanism and processor is extended by speculative execution (in the context of branch prediction) and by a precise interrupt mechanism.

1.1 Overview

The remaining part of this thesis is organized in five chapters:

- Chapter 2 develops the formal framework of our processor. It contains formal approaches and proofs of correctness of a superscalar Tomasulo algorithm, a precise roll-back mechanism and a (speculative) instruction fetch mechanism.
- Chapter 3 describes the implementation of the hardware of our processor design hierarchically descending down to gate-level. The hardware model used for the description is that of [MP95]. It provides a formal background and is easily evaluated.
- Chapter 4 examines the parameter space of the proposed processor design and compares variants of it to existing processor designs from [MP95, MP00, Krö99].
- Chapter 5 defines modules used in the machine description from chapter 3. These modules represent integral functional parts of the processor. Therefore their correctness is of great importance to the correctness of the whole processor design; formal criteria for correctness are specified and proven.

- Chapter 6 draws conclusions from the presented design, the formal approach and its evaluation. It also points out the fields that may be of interest for further research.

1.2 Notes on Reading

We assume that the reader is familiar with the concepts of circuit and processor design as developed in [KP95, MP95, MP00]. For the understanding of a non-superscalar DLX design implementing the Tomasulo algorithm, [Krö99] is of further help.

Chapter 2

Abstract Model

This chapter develops an abstract model for a superscalar hardware in three steps:

- Section 2.1 presents a *superscalar Tomasulo scheduling algorithm*. The algorithm handles the execution of *instruction streams*, sequences of instructions without control flow changes. This limitation is removed later on. The correctness of the algorithm is proven.
- Section 2.2 extends the Tomasulo algorithm by a precise roll-back mechanism. A precise roll-back mechanism stops the execution of an instruction stream at a specified instruction I_n ; the instruction following I_n must not modify the state of the machine according to its sequential semantics. Precise roll-back is needed for the implementation of precise interrupts and speculative execution. The correctness of the roll-back mechanism is proven.
- Section 2.3 develops the concepts of an *instruction fetch mechanism* and of a *speculative instruction fetch mechanism*. These concepts are needed to provide the transition from instruction-stream based machines (that nobody builds) to real-life instruction-memory-based machines (with speculative execution). The speculation makes use of the precise roll-back mechanism to take back the effects of falsely executed instructions.

This chapter does not treat data memory and interrupts specifically. The exact manner of their treatment is often only defined by concrete system architectures. This is not the level of generalization we aim for in this chapter; nevertheless, in chapter 3 data memory and interrupts are treated in the environment of the DLX architecture.

2.1 Tomasulo Algorithm

The section describes a superscalar Tomasulo scheduling algorithm. It executes a dynamic instruction stream I_1, I_2, \dots out-of-order while preserving the sequential semantics.

Our approach is in three steps. First, we define the semantics of an *instruction stream* by a sequential instruction-stream-based machine IS_{seq} . Second, a superscalar Tomasulo machine TM is defined by description of the global structure and the scheduling protocols. In the third step, it is shown that TM simulates IS_{seq} . This is done by showing the data consistency theorem and the termination lemma.

2.1.1 Sequential Instruction-Stream-Based Machine

We define the model of an instruction-stream-based sequential machine, IS_{seq} . The machine IS_{seq} has $\#reg$ registers $R_1, \dots, R_{\#reg}$ over the finite domain DOM . These registers are referenced by the indices $\mathcal{R} := \{1, \dots, \#reg\}$. Instructions are executed from the *instruction stream*. In cycle n the machine executes l_n , a tuple:

$$l_n = (\text{op}, \delta, \text{dop}_1.A, \dots, \text{dop}_\delta.A, \sigma, \text{sop}_1.A, \dots, \text{sop}_\sigma.A)$$

The function $\text{op} : DOM^\sigma \rightarrow DOM^\delta$ computes the results of the operation; δ is the number of destination operands and $\text{dop}_1.A, \dots, \text{dop}_\delta.A \in \mathcal{R}$ are their identifiers (pairwise distinct); σ is the number of source operands and $\text{sop}_1.A, \dots, \text{sop}_\sigma.A \in \mathcal{R}$ are their identifiers. With

$$(\text{result}_1, \dots, \text{result}_\delta) = \text{op}(R_{\text{sop}_1.A}, \dots, R_{\text{sop}_\sigma.A})$$

and primed registers denoting “new”, i.e. next-cycle values of registers, the semantics of the instruction l_n is defined as

$$R'_A := \begin{cases} \text{result}_1 & \text{if } A = \text{dop}_1.A \\ \text{result}_2 & \text{if } A = \text{dop}_2.A \\ \vdots & \vdots \\ \text{result}_\delta & \text{if } A = \text{dop}_\delta.A \\ R_A & \text{otherwise} \end{cases}.$$

Note, that a generalization on multiple, orthogonal destination operands ($\delta > 1$) is not used in real machines. Usually just one general destination operand and additional fixed-addressed destination operands (like operation flags) are sufficient. However, we allow $\delta \geq 1$ for two reasons: first, our approach to fetch mechanisms takes advantage of this fact, justifying the extra effort. Second, the simplifications used in machines having additional fixed-address destinations can be justified on the basis of algorithms for $\delta > 1$.

For ease notation, we consider σ and δ fixed for the instruction set of the machine.

A configuration of the machine stores (the values of) the registers. Instruction l_n is supposed to arrive in cycle n and is therefore not stored in the configuration.

2.1.2 Tomasulo Machine

Informal description of the Tomasulo algorithm

This algorithm was originally published by Tomasulo in [Tom67] in the year 1967. Written first for a very specific environment, this algorithm can be easily adopted for more general architectures. The algorithm associates each instruction and its destination registers on execution with a state-unique identifier, a small natural number, called tag.¹ On requesting a source register, an instruction either receives the (correct) value of this register, or the tag of a previous instruction computing it. On receiving a tag, the instruction has to wait until the result becomes available by a global result broadcast system. If all operands are gathered, an instruction may start execution. On completing execution, the instruction broadcasts tag and results in the whole machine for the benefit of instructions awaiting their source operands. Note that instructions awaiting execution need only to be connected to the global result broadcast system without taking up any other machine resources. A data structure, called reservation station, serves this purpose for a single instruction.

¹State-unique means that in each cycle there is at most one instruction in execution holding the identifier. In the machine, the tags are *recycled* if not used; so the same tag may identify different instructions in different states.

Basic Data Structures and Paths of a Tomasulo Machine

The figure 2.1 shows the basic data structures and paths of a superscalar Tomasulo machine. The figures are annotated with the scheduling phases of an instruction.

The following components are present in a Tomasulo machine:

- The *instruction window* buffers incoming instruction in slots i_1 to i_β . Each round, the machine tries to start the execution of as many of these instructions as possible. According to the definition of the abstract machine, i_i is a tuple of source operand addresses $i_i.sop_1$ to $i_i.sop_\sigma$, an operation code $i_i.op$ and destination operand addresses $i_i.dop_1$ to $i_i.dop_\delta$.

- The *register file* contains a tuple (*valid*, *tag*, *st*, *data*) for each register R_i . Two invariants will hold for the register file:

If $R_i.valid = 1$ then $R_i.data$ contains the data of the last instruction writing to R_i up to the current machine cycle.

If $R_i.valid = 0$, then $R_i.tag$ holds the tag of the *producing instruction*, i.e. the newest instruction having R_i as destination operand. The item $st \in \{1, \dots, \delta\}$ then specifies the index of the destination operand having address i .

- The *common data busses* are the global result broadcast system for the machine. Each common data bus bears a tuple (*tag*, *result*₁, ..., *result* _{δ}) of a tag and the associated results.

- The *reorder buffer* contains for each instruction currently in execution a record, addressed by the instruction's tag, with the following items: the items *dop*₁.A to *dop* _{δ} .A buffer the destination addresses of the instruction. The *valid* flag signals 1, if the instruction has already broadcast its results on a common data bus. In this case, the items *dop*₁.data to *dop* _{δ} .data store the results. The reorder buffer writes results back to the register file in program order.

The reorder buffer is organized as a simple wrap-around queue. The variable *ROB.head* points to the head of the queue, which is also the oldest instruction. The variable *ROB.tail* points to the next free entry of the queue, unless *ROB.tail* = *ROB.head* which signals a full reorder buffer. The constant *ROB.SIZE* denotes the size of the reorder buffer.

- *Reservation station queues* are collections of buffers, called *reservation stations*, for instructions waiting for their source operands to be broadcast on the common data busses. A reservation station contains the following items: the item *full* indicates valid reservation station contents; the item *op* is an operation identifier; the tag of the instruction that the reservation holds is stored in the item *tag*; a tuple *sop* _{σ} for each source operand.

The source operand tuple contains the correct operand data, if *valid* = 1; otherwise it holds the tag and subtag of the producing instruction.

An instruction i_i is executed from the machine in the following way. On *issue* it is taken from the instruction window, associated with a tag from which the reorder buffer can reconstruct program order and put it in an appropriate reservation station. Its source operands may either contain valid data or correct tag and subtag information. The instruction will remain in the reservation station until it has gathered all its missing source operands from the common data busses in a process called *snooping*. Having valid source operand data, the instruction will eventually leave its reservation station and be *dispatched* to its functional unit for *execution*. The functional unit computes the result busses and passes them via an interface called *producer* to one of the common data busses. This step is called *completion*. The

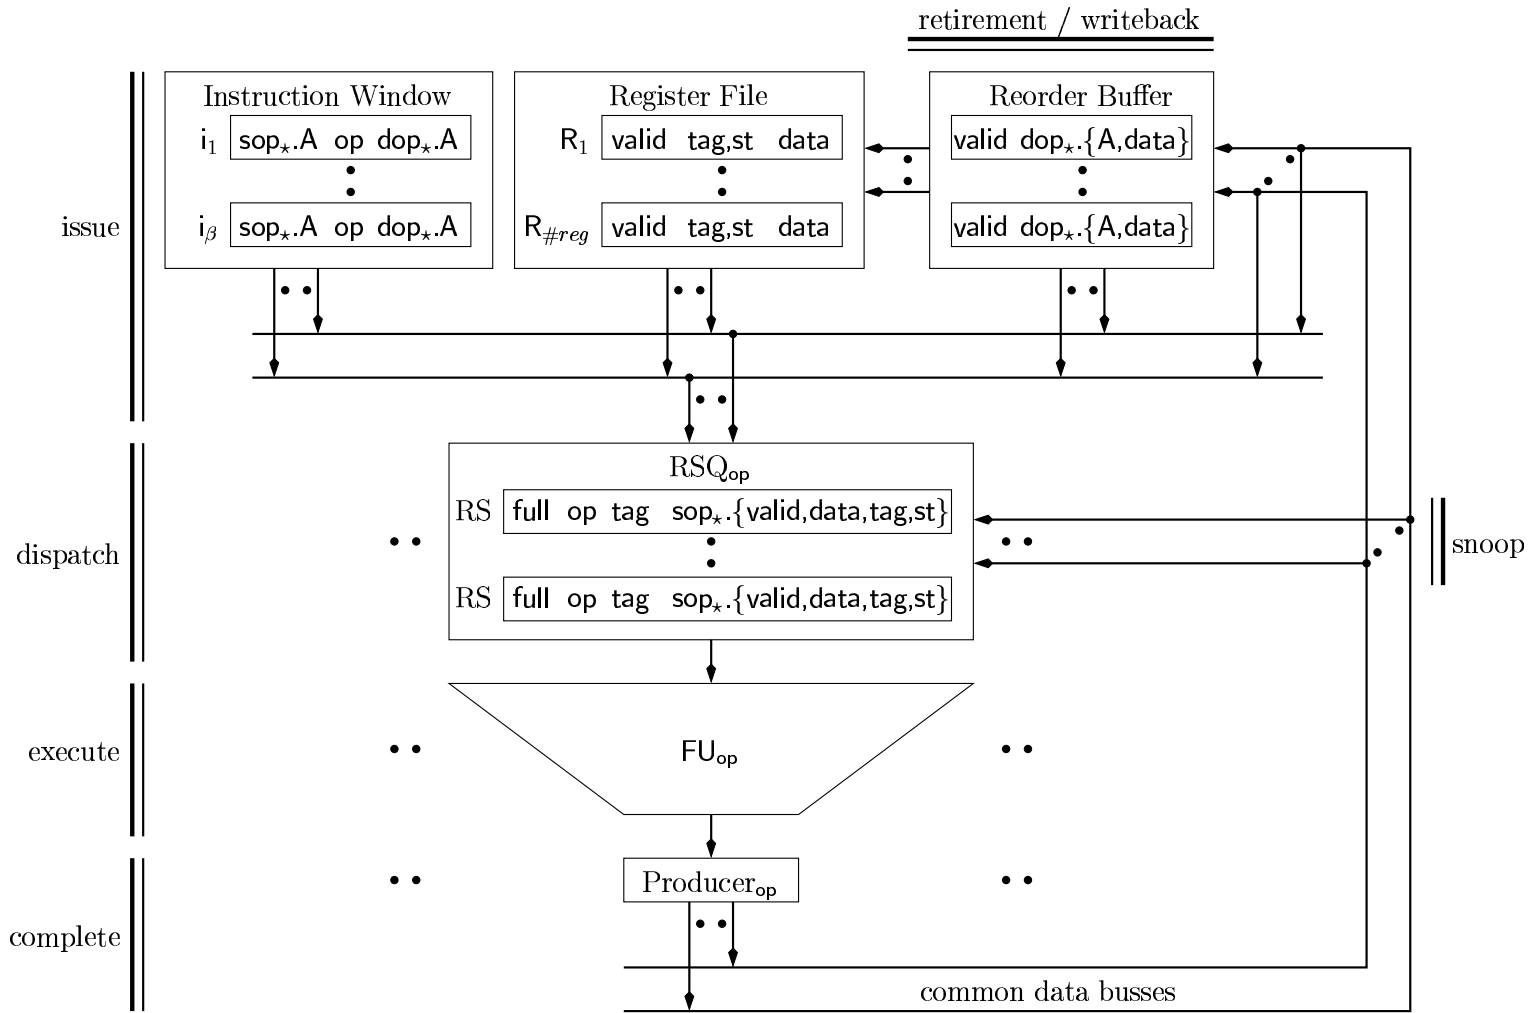


Figure 2.1: Basic data structures and paths of a Tomasulo machine

reorder buffer stores the result and will write it back to the register file in *program order*, i.e. after the results of all previous instructions have been written to the register file. This phase is called *retirement* of an instruction.

The exact procedure of the five phases issue, dispatch, snooping, completion and retirement is defined in protocols below. Basically, these protocols can be obtained by unrolling the non-superscalar Tomasulo algorithm handling only one instruction per cycle.

Issue

Superscalar issuing (algorithm 1) processes the instructions in the instruction window i_1, \dots, i_β with $\beta \in \mathbb{N}$, fixed, in a loop (l. 1) in program order. In line 2, a tag for the instruction i_i is obtained by adding index i to the tail pointer of the ROB. If the tag is not equal to the head pointer (i.e. the ROB is not full) and the appropriate reservation station queue for the instruction is not full, i_i may issue. The test of this condition is contained in line 3.

Instruction issue is divided in three phases: gathering the source operands (ll. 4–27), allocation of a free reservation station (ll. 28–31) and updates for the destination registers in the ROB and the register file (ll. 32–39).

Consider an iteration of the loop body of the source operand construction, ll. 4–27, and define S and R as in ll. 5–6. S is constructed according to five different cases:

First, S might be produced an accompanying instruction from the instruction window (ll. 7–11). This case is indicated if a destination operand's address of the previous instruction matches $S.A$. Naturally, the data is not available for the operand, and we only set the $S.tag$ to the latest instruction's tag having destination $S.A$. The subtag denotes the index of the destination operand.

Second, the register R might be found valid in the register file (ll. 12–14). In this case we set the operand valid and take the data out of the register file.

Third, the operand might be broadcast on one of the common data busses (ll. 15–18). This is the case if the register's tag $R.tag$ is found on one of the common data busses. The data can be taken from the result bus indicated by $R.st$ and stored in the source operand.

Fourth, the operand's data may reside valid in the reorder buffer, waiting for its retirement (ll. 19–21). In this case, the data can be taken from $ROB_{R.tag}.data_{R.st}$.

With none of the previous cases meeting, the operand is still being computed in a functional unit and only its tag $R.tag$ and subtag $R.st$ is available (ll. 22–25).

The destined reservation station RS is filled with the instruction opcode, the destination tag, and the source operands (ll. 28–32). The notation of line 31 results in the assignment of all the items from all the source operands to all the reservation station source operands.

After gathering the source operands, the instruction has to update the register file and the ROB for each destination operand (ll. 32–39). Each destination register is invalidated, the tag information is set to the instruction's tag, the subtag information is set to the destination's index. The addresses of the destination register are stored in the ROB for further use.

In hardware, this protocol is parallelized to handle all the instructions in the instruction window simultaneously. Note the construction in lines 41–43 to stop issuing at the first instruction encountering a full reorder buffer or a non-available reservation station. On termination of the issue protocol, $\#issued$ contains the number of issued instructions.

For notation we define $\#issued^t$ as the value of $\#issued$ in cycle t . The same

```

1: for  $i \leftarrow 1$  to  $\beta$  do
2:    $i_i.\text{tag} \leftarrow (\text{ROB.tail} + i) \bmod \text{ROBSIZE}$ 
3:   if  $i_i.\text{valid} \wedge (i_i.\text{tag} \neq \text{ROB.head}) \wedge \overline{\text{RSQ}_{i_i.\text{op}}.\text{full}}$  then
4:     for all  $\sigma' \in \{1, \dots, \sigma\}$  do (obtain source operands)
5:       Let  $S$  denote  $i_i.\text{sop}_{\sigma'}$ 
6:       Let  $R$  denote  $R_{S.A}$ 
7:       if  $\exists i^- < i : \exists \delta' \in \{1, \dots, \delta\} : i_i^-.\text{dop}_{\delta'}.A = S.A$  then (case: window)
8:         Let  $i^-$  be maximal with  $\exists \delta' \in \{1, \dots, \delta\} : i_i^-.\text{dop}_{\delta'}.A = S.A$ 
9:          $S.\text{valid} \leftarrow 0$ 
10:         $S.\text{tag} \leftarrow i_i^-.\text{tag}$ 
11:         $S.\text{st} \leftarrow \delta'$ 
12:      else if  $R.\text{valid}$  then (case: register)
13:         $S.\text{valid} \leftarrow 1$ 
14:         $S.\text{data} \leftarrow R.\text{data}$ 
15:      else if  $\exists k : \text{CDB}_k.\text{tag} = R.\text{tag}$  then (case: snoop)
16:        Let  $k$  satisfy  $\text{CDB}_k.\text{tag} = R.\text{tag}$ 
17:         $S.\text{valid} \leftarrow 1$ 
18:         $S.\text{data} \leftarrow \text{CDB}_k.\text{data}_{R.\text{st}}$ 
19:      else if  $\text{ROB}_{R.\text{tag}}.\text{valid}$  then (case: forwarding)
20:         $S.\text{valid} \leftarrow 1$ 
21:         $S.\text{data} \leftarrow \text{ROB}_{R.\text{tag}}.\text{data}_{R.\text{st}}$ 
22:      else (case: tag)
23:         $S.\text{valid} \leftarrow 0$ 
24:         $S.\text{tag} \leftarrow R.\text{tag}$ 
25:         $S.\text{st} \leftarrow R.\text{st}$ 
26:      end if
27:    end for
28:    Let  $RS$  denote an empty reservation station in  $\text{RSQ}_{i_i.\text{op}}$ 
29:     $RS.\text{full}' \leftarrow 1$ 
30:     $RS.\text{tag}' \leftarrow i_i.\text{tag}$ 
31:     $RS.\text{sop}'_{*.\star} \leftarrow i_i.\text{sop}_{*.\star}$ 
32:     $\text{ROB}_{i_i.\text{valid}}.\text{valid} \leftarrow 0$ 
33:    for all  $\delta' \in \{1, \dots, \delta\}$  do (update ROB, RF)
34:      Let  $D$  denote  $i_i.\text{dop}_{\delta'}$ 
35:       $R_{D.A}.\text{valid}' \leftarrow 0$ 
36:       $R_{D.A}.\text{tag}' \leftarrow i_i.\text{tag}$ 
37:       $R_{D.A}.\text{st}' \leftarrow \delta'$ 
38:       $\text{ROB}_{i_i.\text{tag}}.\text{dop}_{\delta'}.A' \leftarrow D.A$ 
39:    end for
40:     $i \leftarrow i + 1$ 
41:  else (could not issue  $i$ -th instruction)
42:     $i \leftarrow i - 1$ 
43:    Break out of loop, stop issuing
44:  end if
45: end for
46:  $\text{ROB.tail}' \leftarrow (\text{ROB.tail} + i) \bmod \text{ROBSIZE}$ 
47:  $\#issued \leftarrow i$ 

```

Algorithm 1: Superscalar issue protocol


```

1: for all reservation stations RS do
2:   for all  $\sigma' \in \{1, \dots, \sigma\}$  do
3:     Let S denote RS.sop $_{\sigma'}$ 
4:     if  $\overline{\text{S.valid}} \wedge \exists k : \text{CDB}_k.\text{tag} = \text{S.tag}$  then           (snoop S on CDBk)
5:       Let k satisfy  $\text{CDB}_k.\text{tag} = \text{S.tag}$ 
6:       S.valid'  $\leftarrow 1$ 
7:       S.data'  $\leftarrow \text{CDB}_k.\text{data}_{\text{S.st}}$ 
8:     end if
9:   end for
10: end for

```

Algorithm 2: Superscalar common data bus snooping

```

1: for all functional units FUop do
2:   if  $\overline{\text{FU}_{\text{op}}.\text{stall}} \wedge \exists j : \text{RS}_j.\text{full} \wedge \forall \sigma' : \text{RS}_{\text{op},j}.\text{sop}_{\sigma'}.\text{valid} = 1$  then
3:     RS  $\leftarrow \text{RS}_{\text{op},j}$  with j fairly selected
4:     FUop.op'  $\leftarrow \text{RS.op}$ 
5:     FUop.tag'  $\leftarrow \text{RS.tag}$ 
6:     for all  $\sigma' \in \{1, \dots, \sigma\}$  do           (copy source operands into FUop)
7:       FUop.sop' $_{\sigma'}$   $\leftarrow \text{RS.sop}_{\sigma'}$ 
8:     end for
9:     RS.full'  $\leftarrow 0$ 
10:   end if
11: end for

```

Algorithm 3: Superscalar dispatch

notation defines

$$\mathcal{IW}^t = (i_1^t, \dots, i_\beta^t)$$

as the instruction window presented to the machine in cycle t . The base b of this instruction window is the index of the first instruction in the window. Naturally the base can be computed by $b = 1 + \sum_{1 \leq t' \leq t} \# \text{issued}^{t'}$ and we have

$$(i_1^t, \dots, i_\beta^t) = (l_b, \dots, l_{b+\beta-1})$$

Common Data Bus Snooping

Algorithm 2 defines the superscalar common data bus snooping protocol. The reservation stations snoop on the common data busses to gather their missing source operands' data. I.e., if a reservation station has a non-valid operand and the operand's tag is seen on a CDB (l. 4), then the data is taken from the result bus denoted by the subtag from this CDB and the operand is validated (ll. 5–7). While scanning multiple CDBs at once for an operand's tag, at most one matching CDB may be found because of uniqueness of tags in a single state.

Dispatch

Algorithm 3 shows the superscalar dispatch protocol. Dispatch of a full reservation station to the appropriate functional unit takes place, if all source operands are valid (l. 2). Fair candidate selection (for example by the age of the entries) is needed to ensure termination of instructions in a finite amount of time (l. 3). The reservation copies its operation code, tag and the source operands in the appropriate functional

```

1: for all functional units  $FU_{op}$  do
2:   if  $FU_{op}$  has result and got CDBk-acknowledgement for next cycle then
3:     Let  $T$  denote  $FU_{op}.tag$ 
4:      $ROB_T.valid' \leftarrow 1$ 
5:      $CDB_k.tag' \leftarrow T$ 
6:     for all  $\delta' \in \{1, \dots, \delta\}$  do
7:        $CDB_k.data_{\delta'} \leftarrow FU_{op}.result_{\delta'}$ 
8:        $ROB.dop_{\delta'}.data' \leftarrow FU_{op}.result_{\delta'}$ 
9:     end for
10:   end if
11: end for

```

Algorithm 4: Superscalar completion protocol

Require: register file updates have lower priority than from issue protocol

```

1: for  $i \leftarrow 1$  to  $\epsilon$  do
2:    $r_i.tag \leftarrow (ROB.head + i) \bmod ROBSIZE$ 
3:   Let  $r_i$  denote  $ROB_{r_i.tag}$ 
4:   if  $(r_i.tag \neq ROB.tail) \wedge r_i.valid$  then
5:     for all  $\delta' \in \{1, \dots, \delta\}$  do (update RF)
6:       Let  $A$  denote  $r_i.dop_{\delta'}.A$ 
7:        $R_A.data' \leftarrow r_i.dop_{\delta'}.data$ 
8:       if  $r_i.tag = R_A.tag$  then (validate register on tag equality)
9:          $R_A.valid' \leftarrow 1$ 
10:      end if
11:    end for
12:  else (no valid instructions left for retirement)
13:     $i \leftarrow i - 1$ 
14:    Break out of loop, stop retiring.
15:  end if
16: end for
17:  $ROB.head' \leftarrow (ROB.head + i) \bmod ROBSIZE$ 
18:  $\#retired \leftarrow i$ 

```

Algorithm 5: Superscalar retirement

unit (ll. 4–8). After dispatch has taken place, the reservation station is marked as empty (l. 9).

Note, that dispatch to duplicates of functional units is only a simple extension to the protocol.

Completion

The completion protocol (algorithm 4) puts instruction results on a common data bus. If a results is available and a CDB has been acknowledged by the CDB bus control (l. 2) the following actions are performed: The associated ROB entry is marked valid in l. 4; the CDB tag field is set in l. 5. The results of the instruction are copied on the CDB and into the ROB (ll. 6–9).

Retirement

During retirement, completed instructions are taken in program order out of the reorder buffer and their results are stored in the register file. Without interrupts,

the reorder buffer is not a necessity. However, reordering instructions is required for the precise roll-back mechanism developed in section 2.2.

Algorithm 5 defines the retirement protocol. It exploits the queue organization of the reorder buffer, with the ROB head pointing to the oldest instruction present in the reorder buffer. The retirement protocol is defined in phases. The i -th phase attempts to retire instruction r_i located at position $\text{ROB.head} + i \bmod \text{ROBSIZE}$ in the queue.

Instruction r_i may retire, if it still is a valid queue entry ($r_i.\text{tag} \neq \text{ROB.tail}$) and $r_i.\text{valid} = 1$, i.e. the instruction has completed (l. 5).

All destination operands are written back into the register file (ll. 6–8). A tag comparison (l. 9) indicates, if the retiring instruction is still the producing instruction for register R . In this case, the register entry is validated (l. 10).

Note that the issue protocol and the retirement protocol concurrently write to the register file. These conflicts are resolved by giving the issue protocol a writing priority over the retirement protocol. The proof of correctness relies on this behaviour.

Like the issue protocol the retirement returns the number of processed instructions in a variable, $\#retired$. Again, we define for notation $\#retired^t$ as the value of $\#retired$ in cycle t . The sequence

$$\mathcal{R}^t = (r_1.\text{tag}^t, \dots, r_{\#retired^t}.\text{tag})$$

denotes the tags of the retiring instructions in cycle t ; $r_i.\text{tag}$ is the tag of the instruction processed in the i -th retirement phase.

2.1.3 Correctness

This section proves the correctness of the superscalar Tomasulo algorithm in two steps. In the first step, data-consistency is shown: all instructions receive the correct source operands as defined by the (sequential) semantics. In the second step the termination of the protocols is examined: instructions take only a finite amount of time to process.

At first we define for notation some *partial* functions:

$$\begin{aligned} \text{last}(A, n) &= \max \{n' < n \mid \exists \delta' : l_{n'}.\text{dop}_{\delta'}.A = A\} \\ \text{st}(A, n) &= \delta' \text{ with } l_n.\text{dop}_{\delta'}.A = A \\ \text{result}(A, n) &= l_n.\text{result}_{\delta'} \text{ with } \delta' = \text{st}(A, n) \\ \text{tag}(A, n) &= l_n.\text{tag} \text{ if } \exists \delta' : l_{n'}.\text{dop}_{\delta'}.A = A \\ \text{idx}(\text{tag}, t) &= \max \{n' \mid l_{n'}.\text{tag} = \text{tag} \wedge l_{n'} \text{ issued in cycle } t' \leq t\} \end{aligned}$$

In words the notations read as follows: The index of the last instruction writing to R_A before instruction l_n is denoted by $\text{last}(A, n)$. The subtag, i.e. index of destination operand, of instruction l_n writing to R_A is denoted by $\text{st}(A, n)$. The result of an instruction's write to R_A is denoted by $\text{result}(A, n)$. The tag that has been given to instruction l_n —if it writes R_A —is denoted by $\text{tag}(A, n)$. Finally, $\text{idx}(\text{tag}, t)$ is the index n of the instruction lastly associated with the tag tag in cycle t .

The notation C^t for some component C denotes the value of this component at the beginning of cycle t .

Theorem 2.1 *Have a reservation station RS with $\text{RS.full}^t = 1$. Let n be the index of the instruction that RS holds, i.e. $n := \text{idx}(\text{RS.tag}^t, t)$. Then for all $1 \leq \sigma' \leq \sigma \wedge A := l_n.\text{sop}_{\sigma'}.A$:*

$$\text{RS.sop}_{\sigma'}.\text{valid}^t = 0 \implies \text{RS.sop}_{\sigma'}.\text{tag}^t = \text{tag}(A, \text{last}(A, n))$$

$$\begin{aligned} \text{RS.sop}_{\sigma'}.st^t &= st(A, \text{last}(A, n)) \\ \text{RS.sop}_{\sigma'}.valid^t = 1 &\implies \text{RS.sop}_{\sigma'}.data^t = \text{result}(A, \text{last}(A, n)) \end{aligned}$$

Corollary 2.2 (Data Consistency) *If the functional units work correctly (and terminate), the superscalar Tomasulo algorithm is data consistent.*

Proof of Corollary 2.2. Once an instruction becomes dispatched, all its source operands are valid. By the theorem, all source operand data fields contain the result of the last instruction writing to them. The functional unit therefore produces the correct result.

Proof of Theorem 2.1. We simultaneously prove the claim and the following invariant:

Invariant 2.3 *Let A be a register address and b the base of the instruction window \mathcal{IW}^t in cycle t . Then the following implications hold:*

$$\begin{aligned} R_A.valid^t = 0 &\implies R_A.tag^t = \text{tag}(A, \text{last}(A, b)) \\ &\quad R_A.st^t = st(A, \text{last}(A, b)) \\ R_A.valid^t = 1 &\implies R_A.data^t = \text{result}(A, \text{last}(A, b)) \end{aligned}$$

The theorem is proved by induction over the cycles of the machine t . For the induction basis we have $t = 1$. Since an initialization condition of the machine is assumed, the reservation stations are empty and the registers are valid and zero-valued; the claims of the theorem and the invariant therefore hold trivially. Now assume that $t > 1$ and the invariant, the theorem and the corollary on data consistency hold for all $t' < t$. Let b be the base of \mathcal{IW}^t , $t^- := t - 1$ and b^- be the base of \mathcal{IW}^{t^-} .

First, the induction step for the invariant is shown. Define n^- maximal with the property $b^- \leq n^- < b \wedge \exists \delta' : l_{n^-}.dop_{\delta'}.A = A$. The following two cases must be examined:

- There was an issuing instruction with destination register R_A in cycle t^- . This is equivalent to n^- having a defined value. Then, obviously $\text{last}(A, b) = n^-$. In cycle t^- , the issue protocol sets

$$\begin{aligned} R_A.valid^t &= 0 \\ R_A.tag^t &= \text{tag}\left(A, b^- + \max\left\{i \leq \#\text{issued}^{t^-} \mid \exists \delta' : i_i.dop_{\delta'}^{t^-} = A\right\}\right) \\ &= \text{tag}(A, n^-) \\ &= \text{tag}(A, \text{last}(A, b)) \\ R_A.st^t &= st\left(A, b^- + \max\left\{i \leq \#\text{issued}^{t^-} \mid \exists \delta' : i_i.dop_{\delta'}^{t^-} = A\right\}\right) \\ &= st(A, n^-) \\ &= st(A, \text{last}(A, b)) \end{aligned}$$

Note that in this case retiring instructions are (correctly) not taken into account, since the issue protocol has a higher write priority than the retirement protocol on the register file (cf. section 2.1.2).

- Now assume that we are not in the first case, i.e. there was no issuing instruction with destination register R_A in cycle t^- . Therefore:

$$\text{last}(A, b) = \text{last}(A, b^-)$$

Further assume, there was a retiring instruction in cycle t^- that matched tag with $R_A.\text{tag}^{t^-}$; i.e.

$$\exists i \leq \#retired^{t^-} : r_i.\text{tag} = R_A.\text{tag}$$

There is at most one such i , because no two retiring instructions can have the same tag. The retirement protocol in cycle t^- correctly sets:

$$\begin{aligned} R_A.\text{valid}^t &= 1 \\ R_A.\text{data}^t &= \text{result}(A, \text{idx}(r.\text{tag})) \\ &= \text{result}(A, \text{last}(A, b^-)) \\ &= \text{result}(A, \text{last}(A, b)) \end{aligned}$$

If there was no retiring instructions matching tag with $R_A.\text{tag}$, the register, as well as $\text{last}(A, b^-)$ remain unchanged and the induction assumption holds.

This shows the induction step for the invariant. For the induction step of the theorem two cases must be distinguished:

- $RS.\text{full}^{t^-} = 0$. With the assumption of the theorem, $RS.\text{full}^t = 1$, this means that an issue took place on the reservation station in cycle t^- .

If $RS.\text{sop}_{\sigma'}.\text{valid}^t = 1$ the issue protocol has filled $RS.\text{sop}_{\sigma'}.\text{data}$ from the register file, from a common data bus CDB_k or from the reorder buffer. Because of the induction assumption on the invariant and on data consistency, each case results in:

$$\begin{aligned} RS.\text{sop}_{\sigma'}.\text{data} &= \begin{cases} CDB_k.\text{dop}_{\delta'}.\text{result} \\ ROB_{R_A.\text{tag}}.\text{data} \\ R_A.\text{data} \end{cases} \\ &= \text{result}(A, \text{last}(A, n)) \end{aligned}$$

Otherwise, assume $RS.\text{sop}_{\sigma'}.\text{valid}^t = 0$. If there is an n^- , maximal, with the property $b^- \leq n^- < n \wedge \exists \delta' : l_{n'}.\text{dop}_{\delta'}.\text{A} = A$ the tag and the subtag were set by tag forwarding:

$$\begin{aligned} RS.\text{sop}_{\delta'}.\text{tag} &= \text{tag}(A, b^- + \max \{i \leq \#issued^{t^-} \mid \exists \delta' : i_i.\text{dop}_{\delta'}^{t^-} = A\}) \\ &= \text{tag}(A, n^-) \\ &= \text{tag}(A, \text{last}(A, n)) \\ RS.\text{sop}_{\delta'}.\text{st} &= \text{st}(A, b^- + \max \{i \leq \#issued^{t^-} \mid \exists \delta' : i_i.\text{dop}_{\delta'}^{t^-} = A\}) \\ &= \text{st}(A, n^-) \\ &= \text{st}(A, \text{last}(A, n)) \end{aligned}$$

If there is no defined n^- , then then tag and subtag are taken out of the register file. By the induction assumption for the invariant:

$$\begin{aligned} RS.\text{sop}_{\delta'}.\text{tag} &= R_A.\text{tag}^{t^-} \\ &= \text{tag}(\text{last}(A, b^-)) \\ &= \text{tag}(A, \text{last}(A, n)) \\ RS.\text{sop}_{\delta'}.\text{st} &= R_A.\text{st}^{t^-} \\ &= \text{st}(\text{last}(A, b^-)) \\ &= \text{st}(A, \text{last}(A, n)) \end{aligned}$$

- $\text{RS.full}^{t^-} = 1$. The reservation station was already full in the last cycle. If $\text{RS.sop}_{\sigma'}.valid^t = \text{RS.sop}_{\sigma'}.valid^{t^-}$, source operand $\text{RS.sop}_{\sigma'}$ did not change its value in the preceding cycle and the claim holds by induction assumption.

Otherwise, have $\text{RS.sop}_{\sigma'}.valid^{t^-} = 0 \wedge \text{RS.sop}_{\sigma'}.valid^t = 1$. This means that the reservation station snooped operand $\text{RS.sop}_{\sigma'}$ in the last cycle on, say, CDB_j . By induction assumption, have

$$\begin{aligned} \text{CDB}_j.\text{tag}^{t^-} &= \text{RS.sop}_{\sigma'}.\text{tag}^{t^-} \\ &= \text{tag}(A, \text{last}(A, n)) \end{aligned}$$

which implicates by the induction assumption on data consistency:

$$\begin{aligned} \text{RS.sop}_{\sigma'}.\text{data}^t &= \text{CDB}_j.\text{result}_{\text{RS.sop}_{\sigma'}.st}^{t^-} \\ &= \text{result}(A, \text{last}(A, n)) \end{aligned}$$

Lemma 2.4 (Termination) *If CDB acknowledgements are fair and computations in functional units take up finite time, then the Tomasulo algorithm does not deadlock the machine.*

Proof. This lemma is also proven by induction over the instructions in the instruction stream.

We show that the first instruction terminates. For the first instruction, all register file items are valid, the reorder buffer and all the reservation stations are empty. Since the first instruction is also first in the instruction window, the instruction window case does not apply. Therefore, l_0 is issued to the appropriate reservation station with all operands valid according to the register case in the issue protocol. Dispatch takes place one cycle after issue. Since the functional unit takes only finite time to compute the result of the instruction and because of fair CDB acknowledgement, completion takes place. Retirement follows one cycle after completion, since l_0 occupies the ROB head.

Now assume that instructions l_0, \dots, l_{i-1} terminate. Issuing takes place if room is available in the ROB and the appropriate reservation station queue. Eventually, this is the case: according to the induction assumption l_0, \dots, l_{i-1} terminate and especially leave the reservation stations and the ROB. Having space, the issue protocol processes l_i . The issue protocol states that each instruction looks for its source operands at all places where results can be found. If a source operand is not valid immediately, it is snooped on the common data busses. Therefore and because of fair dispatch acknowledgements, the instruction dispatches. The functional unit executing l_i only takes finite time to compute the result of the instruction. Because of fair CDB acknowledgements, the instruction is completed in finite time. Retirement also takes finite time, since by induction assumption all the previous instructions retire.

2.2 Roll-Back

This section shows the implementation of a precise roll-back mechanism. A roll-back mechanism allows to suspend the execution of an instruction stream at a retiring instruction having $\text{tag } r_i.\text{tag} \in \mathcal{R}^t$, a non-trivial operation on pipelined or superscalar machines. If $n := \text{idx}(r_i.\text{tag}, t)$, preciseness requires that after roll-back, the instructions l_1, \dots, l_n have been retired and the instructions l_{n+1}, l_{n+2}, \dots do not have any effect (this definition reminds of the definition of precise interrupts [Mül97, SP88]; it is formulated differently below).

A precise roll-back mechanism is an essential for support of precise interrupts and speculative execution. *Interrupts* force the machine to switch to the interrupt service routine's instruction stream on internal or external interrupt conditions. Preciseness of interrupts requires this switch to be made in well-defined and recoverable manner. *Speculative execution* needs to discard the instructions of a wrongly predicted instruction stream and to continue execution with the correct one.

This section proceeds in three steps. First, the formal definition of a precise roll-back for a superscalar machine is given. Second, precise roll-back support is incorporated in the retirement protocol. Third, we prove that the new protocol meets the formal requirements of a precise roll-back.

2.2.1 Definition of Precise Roll-Back

The semantics of a precise roll-back are defined with the help of two invariants on the semantics of a sequential machine.

Definition 2.5 (Sequential Configuration) *A configuration for the sequential machine IS_{seq} defined in section 2.1 is a tuple of the values of all its registers and the index n of the lastly executed instruction.*

Definition 2.6 (Projection on Sequential Configuration) *The projection of the Tomasulo machine TM after retirement phase i in cycle t is defined as*

$$\text{pr}_i(C_{TM}^t) = (\text{idx}(r_i.\text{tag}^t, t), \forall A : R_{i,A}.\text{data}')$$

where $R_{i,A}.\text{data}'$ denotes the value of $R_A.\text{data}'$ after the i -th retirement phase.

Definition 2.7 *A machine having a precise roll-back mechanism must fulfil the following two invariants:*

- The projection invariant states that in each retirement phase, the register file contains the correct register data values for the instruction up to the retiring instruction:

$$\forall 1 \leq i \leq \#retired^t, n := \text{idx}(r_i.\text{tag}, t) : \text{pr}_i(C_{TM}^t) = C_{seq}^n$$

- The halting invariant states that on executing a roll-back after instruction l_n in cycle t and after retirement phase i ($n := \text{idx}(r_i.\text{tag}, t)$) must not change the projected configuration on empty instruction stream:

$$\forall t' > t, i_1.\text{valid}^{t'} = 0 : \forall 1 \leq i' \leq \#retired^{t'} : \text{pr}_{i'}(C_{TM}^{t'}) = C_{seq}^n$$

Remark 2.8 *The projection invariant ensures the projected state of the machine equal to the state of the sequential machine at defined times. The halting variant implicitly ensures that the machine does not do something “nasty” while it has been told to roll-back. As can be seen below, this criterion is easily verified for the Tomasulo algorithm.*

2.2.2 Roll-Back Retirement Protocol

To implement a precise roll-back mechanism, the retirement protocol, algorithm 5, has to be modified. Algorithm 6 shows the extensions, which may be inserted between line 12 and line 13 of the original algorithm 5.

The method implemented in our machine is called roll-back by *flushing*. The machine simply wipes away all the information of instructions still resident in the machine by clearing all RSQs, the ROB and the computations in functional units (ll. 2–8).

We call the machine using the extended retirement protocol TM_{rb} .

- 1: **if** perform roll-back? **then**
- 2: Update as follows with priority over the other protocols:
- 3: Set $\text{RS.full}' \leftarrow 0$ for all reservation stations
- 4: Clear all computations in all functional units FU_{op}
- 5: $\forall A : \text{R}_A.\text{valid}' \leftarrow 1$
- 6: $\text{ROB.head}' \leftarrow 0$
- 7: $\text{ROB.tail}' \leftarrow 0$
- 8: Break out of loop (do not modify head)
- 9: **end if**

Algorithm 6: Extension to superscalar retirement to support precise roll-back

2.2.3 Proof of Preciseness

Lemma 2.9 *The machine TM_{rb} satisfies the projection invariant.*

We proof the projection invariant in the following form:

$$\forall t, 1 \leq i \leq \#retired^t, n := \text{idx}(r_i.\text{tag}, t), A : \text{R}_A.\text{data}' = \text{result}(A, \text{last}(A, n))$$

In words: in every retirement phase i each register's data contains the result of the last instruction writing to it.

The proof is by induction over the cycles t . Let $1 \leq i \leq \#retired^t$, $n := \text{idx}(r_i.\text{tag}, t)$. For $t = 1$ there is nothing to show, since no instructions retires and TM_{rb} and IS_{seq} are assumed to be initialized in the same way.

Let $t > 1$ and consider the claim true for any $t' < t$. Let A be a register address. If no instructions prior to and including the i -th retirements phase writes to R_A , the claim holds because of the induction assumption.

Otherwise, $\text{R}_A.\text{data}'$ contains the data of the latest retirement writing to R_A . By data consistency and the queue order of the ROB, this is also the data of the last instruction writing R_A , so

$$\text{R}_A.\text{data}' = \text{result}(A, \text{last}(A, n))$$

Lemma 2.10 *The machine TM_{rb} satisfies the halting invariant.*

Proof. This lemma is proven by taking a quick look at all the protocols. The issue protocol terminates directly on seeing $i_1.\text{valid} = 0$ and does not change ROB.tail . The snooping protocol and the dispatch protocol only operate on full reservation stations of which there are none. The functional units have been told to stop any computations, so the completion protocols does not find a result to forward to the reorder buffer. The retirement protocol terminates directly on $\text{ROB.head} = \text{ROB.tail}$ without changing ROB.head .

2.3 Instruction Fetch and Speculation

2.3.1 Sequential Instruction-Memory-Based Machine

We define a sequential, instruction-memory-based machine, IM_{seq} . The definition of IM_{seq} is similar to the definition of IS_{seq} given in section 2.1.1. The machine IM_{seq} also has $\#reg$ registers $\text{R}_1, \dots, \text{R}_{\#reg}$ over the finite domain DOM . Additionally, the machine has a special register, called the program counter register PC .

The $\text{PC} \in \{1, \dots, N\} \subseteq \text{DOM}$ for some $N \in \mathbb{N}$ determines the instructions that IM_{seq} executes from the instruction memory $\mathcal{IM} = (\text{IM}_1, \dots, \text{IM}_N)$. In the n -th cycle IM_{seq} executes the instruction $I_n := \text{IM}_{\text{PC}^n}$ with R^n denoting the value of

the register R in the n -th cycle. The sequence of instructions I_1, I_2, \dots is called the dynamic instruction stream, while the instructions IM_i stored in the instruction memory are called the static instructions. The instructions IM_i in the instruction memory are tuples of the following form:

$$IM_i = (\text{op}, \delta, \text{dop}_1.A, \dots, \text{dop}_\delta.A, \sigma, \text{sop}_1.A, \dots, \text{sop}_\sigma.A)$$

If the set of register identifiers is extended to $\mathcal{R} = \{1, \dots, \#reg\} \cup \{PC\}$, the tuples are of the same type as in the instruction-stream-based machine.

The semantics of the regular registers is identical to the instruction-stream-based machine. With

$$(\text{result}_1, \dots, \text{result}_\delta) = \text{op}(R_{\text{sop}_1.A}, \dots, R_{\text{sop}_\sigma.A})$$

the program counter is updated in the n -th cycle according to the following equations:

$$PC' := \begin{cases} \text{result}_1 & \text{if } \text{dop}_1.A = PC \\ \text{result}_2 & \text{if } \text{dop}_2.A = PC \\ \vdots & \vdots \\ \text{result}_\delta & \text{if } \text{dop}_\delta.A = PC \\ PC + 1 & \text{otherwise} \end{cases}$$

As an initialization condition, we set $PC^1 = 1$.

2.3.2 Non-Speculative Instruction Fetch

This section develops the notion of an instruction fetch mechanism. Informally speaking, an instruction fetch mechanism is a machine that constructs the dynamic instruction stream out of the static instructions. A machine simulating IM_{seq} naturally constructs the dynamic instruction stream at some time and may therefore be called an instruction fetch mechanism for IM_{seq} . What we are rather interested in are simple machines providing the dynamic instruction stream for execution on a Tomasulo machine.

Definition

We construct the Tomasulo machine TM' , an extension of TM . The following three modifications are made:

- The instruction window's entries are extended by an additional item PC containing the program counter of the associated instruction. This information can be used by the machine, whenever PC is needed as a source operand. So, an instruction window entry i_i is a tuple

$$i_i = (\text{valid}, PC, \text{op}, \delta, \text{dop}_1.A, \dots, \text{dop}_\delta.A, \sigma, \text{sop}_1.A, \dots, \text{sop}_\sigma.A)$$

- The issue protocol allows read access to its variable $\#issued$.
- The retirement protocol is modified to report all write accesses to the PC to an external interface. This modification allows for simple fetch mechanism, as the results of computations of the Tomasulo machine can be used.

The necessary protocol changes are trivial and not included here for the sake of brevity.

Now an instruction fetch mechanism is a machine that correctly constructs the dynamic instruction stream for execution with our Tomasulo machine TM' and without deadlock:

```

1: loop
2:   Fetch a basic block starting at fetchPC
3:   Wait for the CFI to execute
4:   Set fetchPC to the reported result
5: end loop

```

Algorithm 7: Generic scalar instruction fetch mechanism

Definition 2.11 (Instruction Fetch Mechanism) *A machine IFM is called an instruction fetch mechanism for TM' if it constructs instruction windows \mathcal{IW}^t satisfying the following property:*

For every n there is an index t such that the sequence of issued instruction and their PC of TM' in cooperation with IFM is a prefix of the sequence of issued instructions and their PC of IM_{seq} .

In conclusion we find the following theorem:

Theorem 2.12 *An instruction fetch mechanism IFM and TM' in cooperation simulate IM_{seq} .*

Scalar Fetch Mechanism

We develop a simple fetch mechanism, called *scalar fetching* (cf. [Joh91]). The instructions are divided into two classes, the basic-block instructions (BBI) and the control flow instructions (CFI):

$$\begin{aligned}
IM_i \text{ is BBI} & : \iff \forall 1 \leq \delta' \leq \delta : IM_i.\text{dop}_{\delta'} \neq PC \\
IM_i \text{ is CFI} & : \iff \exists 1 \leq \delta' \leq \delta : IM_i.\text{dop}_{\delta'} = PC
\end{aligned}$$

Definition 2.13 *Let $i, l \in \mathbb{N}$. The finite instruction sequence $(IM_i, \dots, IM_{i+l-1})$ with*

$$IM_{i+l} \text{ is CFI} \quad \wedge \quad \forall j \in \{i, \dots, i+l-1\} : IM_j \text{ is BBI}$$

is called the basic block starting at i with length l . If $l = 0$ then the basic block starting at i is said to be empty.

The following lemma contains an elementary but significant observation about the semantics of our new machine:

Lemma 2.14 *Let $I_t = IM_{PC^t}$. Let l be the length of the basic block starting at PC. Then:*

$$\forall l' \in \{0, \dots, l\} : I_{n+l'} = IM_{PC^t+l'}$$

Proof. BBIs inside a basic block do not modify the PC, so execution is sequential.

This way, the dynamic instruction sequence $I = (I_0, I_1, \dots)$ can be rewritten as an alternating sequence of basic blocks and CFIs:

$$I = (bb_0, cfi_0, bb_1, cfi_1, \dots)$$

In this equation, bb_i is the basic block starting at $PC(cfi_{i-1})$ and cfi_i is the CFI following bb_{i-1} . The basic blocks are possibly empty.

Lemma 2.14 justifies the generic scalar fetch mechanism for TM' shown in Algorithm 7. The fetch mechanism fetches a blocks of BBIs followed by a CFI in the

first step (line 2). After that it waits for the machine to execute the CFI (line 3). According to the communication protocol of the fetch mechanism and the machine, the machine notifies the fetch mechanism of the newly computed value for PC. This is used by the fetch mechanism to update its `fetchPC` in the third step (line 3). Traditional and especially non-superscalar fetch mechanisms are based on this algorithm.

2.3.3 Speculative Instruction Fetch

Scalar instruction fetch much too slow for superscalar machines. This is due to the small average basic block length of about 5 instructions. This means, while superscalar machines intend to process many instructions in parallel, they have wait for new instructions every few cycles due to scalar instruction fetch.

Today's microprocessors have the ability to execute instructions *speculatively*. This means that for each CFI they can guess a target PC prior to its computation. The speculation must be verified in the Tomasulo machine. If verification is successful, the machine may continue execution; otherwise the machine must roll back and proceed execution on the "right" execution path.

Definition

We proceed the same way as for non-speculative instruction fetch. First, we define a Tomasulo machine TM_{spec} supporting speculative execution. Then we define speculative instruction fetch mechanisms for TM_{spec} .

The machine TM_{spec} is based on TM_{rb} and TM' with the following modifications:

- The instruction window's entries are extended by additional entries `bbi`, `cfi` and `nPC`. The boolean variable `bbi` indicates if the fetch mechanism "believes" the associated instruction to be a basic-block instruction. The boolean variable `cfi` indicates if the fetch mechanism "believes" the associated instruction to be a control flow instruction. The entry `nPC` indicates the alleged PC after the execution of the instruction, i.e. the source PC for the following instruction. So, an instruction window entry i_i is a tuple

$$i_i = (\text{valid}, \text{PC}, \text{nPC}, \text{bbi}, \text{cfi}, \\ \text{op}, \delta, \text{dop}_1.A, \dots, \text{dop}_\delta.A, \sigma, \text{sop}_1.A, \dots, \text{sop}_\sigma.A)$$

- The issue protocol is modified to stop issuing on encountering an instruction that is not properly marked as BBI or CFI instruction. The fetch mechanism is notified of that condition. Furthermore, variable `#issued` is passed back to the fetch mechanism.
- The retirement protocol checks for every retiring instruction, if the speculated value of the next PC, `nPC`, matches the computed value of the next PC. If this is not the case, the machine notifies the fetch mechanism and performs a roll-back immediately *after* the offending instruction.
- Again, the retirement protocol is modified to report all write accesses to the PC to an external interface.

Algorithms 8 and 9 list the modifications necessary to the old algorithms 5 and 6. The remaining modifications again are considered straightforward and not listed here for the sake of brevity.

As can be seen, a speculative instruction fetch involves two levels of speculation.

```

1: if  $i_i$ .valid and  $i_i$  has missing or wrong mark in CFI / BBI field then
2:   Notify fetch mechanism.
3:    $i \leftarrow i - 1$ 
4:   Break out of loop, stop issuing
5: end if

```

Algorithm 8: Addition to the issue protocol

```

1: if  $\exists \delta' \in \{1, \dots, \delta\} \wedge r_i.\text{dop}_{\delta'}.A = \text{PC}$  then
2:   if  $r_i.\text{dop}_{\delta'}.result \neq \text{ROB}_B.\text{nPC}$  then
3:     Initiate roll-back. Notify fetch mechanism.
4:   else
5:     Notify fetch mechanism of verification.
6:   end if
7: end if

```

Algorithm 9: Addition to the retirement protocol

```

1: loop
2:   if retirement protocol detects misspeculation then
3:     Set PC to the corrected version
4:     Invalidate the instruction window
5:   else
6:     Fetch a basic-block
7:     Guess a target for the CFI and modify PC
8:     Annotate the instructions
9:   end if
10: end loop

```

Algorithm 10: Generic speculative instruction fetch mechanism

First, the instruction fetch mechanism is allowed to speculate about the class of an instruction. This approach will not be followed in this thesis, although it is common for high performance fetch mechanism. For example, instruction fetch mechanisms might cache the class-information of instructions basicblock-wise, to avoid the decoding of instructions (cf. [Yeh93]).

Second, it speculates control flows, i.e. the irregular change of the PC.

Both speculations are verified and may cause a roll-back. In the first case, this is detected as early as in the issue phase. As the offending instruction is kept from issuing, no “real” roll-back is necessary. In the second case, the verification takes place in the retirement protocol. A misspeculation causes a roll-back taking back the effects of the instructions still in execution.

Now we define:

Definition 2.15 (Speculative Instruction Fetch) *A machine SIFM is called an instruction fetch mechanism for TM_{spec} if it constructs instruction windows IW^t so that the following property holds:*

For every n there is an index t such that the sequence of retired instruction and their PC of TM_{spec} in cooperation with SIFM is a prefix of the sequence of issued instructions and their PC of IM_{seq} .

A Generic Speculative Instruction Fetch Mechanism

```

1: if rollback to cfcPC signalled then
2:   Remove all instruction fetch queues
3:   Let  $Q$  be a new queue
4:   fetchIFQ  $\leftarrow Q$ 
5:   issueIFQ  $\leftarrow Q$ 
6:   fetchPC  $\leftarrow$  cfcPC
7: end if
8:  $i \leftarrow 0$ 
9: for  $i \leftarrow 1$  to  $\alpha$  do
10:  if  $IM_{\text{fetchPC}}$  is CFI then
11:    Guess a result cfcPC for the update of the fetchPC register
12:    fetchIFQ.push ("valid CFI", PC = fetchPC, nPC = cfcPC)
13:    fetchPC  $\leftarrow$  cfcPC
14:    fetchIFQ  $\leftarrow$  new queue  $Q$ 
15:    Break out of loop. Stop fetching.
16:  else
17:    fetchIFQ.push ("valid BBI", PC = fetchPC, nPC = fetchPC + 1)
18:    fetchPC  $\leftarrow$  fetchPC + 1
19:  end if
20: end for
21:  $i_{*} \leftarrow$  issueIFQ.*
22: issueIFQ.drain (#issued)
23: if issueIFQ.empty  $\wedge$  (fetchIFQ  $\neq$  issueIFQ) then
24:   issueIFQ  $\leftarrow$  next IFQ
25: end if

```

Algorithm 11: Queue-based speculative instruction fetch

Algorithm 10 shows the algorithm for a generic speculative instruction fetch mechanism. If the retirement protocol detects a misspeculation, it invalidates the instruction window and sets its internal PC to the corrected version (ll. 2–4). Otherwise it fetches a basic-block. Each basic-block ends in a CFI. The outcome the CFI is guessed and the PC is modified accordingly (ll. 6–8). The process repeats.

Queue-Based speculative instruction fetch

The generic speculative instruction is not of much use in real implementations. We present here a different algorithm closer to the implementation of chapter 3.

The algorithm maintains a list of queues, each of which will store a single basic block. The instruction fetch is decoupled from the instruction issue: it addresses the memory via the fetch PC, fetchPC, and appends instructions to a designated fetching queue until a CFI has been met. The PC computation of a CFI will be guessed and fetching continues in a new queue from the guess control flow change PC, cfcPC.

Instructions will be issued by draining them from a designated issuing IFQ. If the issuing IFQ is empty, issuing switches to the the next queue.

Initialization and rollback result in the clearing of all queues and starting all over with a new queue Q , which will be fetching and issuing IFQ.

Algorithm 11 shows the algorithm just described.

Chapter 3

Hardware

This chapter develops the hardware of a superscalar DLX processor implementing the Tomasulo algorithm. The basic structure of such a processor is in analogy to a non-superscalar version; [Krö99] develops a non-superscalar processor. Figure 3.1 on the following page shows an overview of the data paths of our processor design. The processor decomposes the execution of an instruction into five pipeline stages. In the first, the instruction fetch mechanism accesses the instruction memory, predicts control flow and puts instructions in the instruction window. From the instruction window, instructions are decoded and led on to the appropriate reservation station in the second stage. During decode and in the reservation stations the instructions gather their source operands, which may come out of the register file, the reorder buffer or from a common data bus. Instructions with available source operands are led to the third stage, the execution stage. This stage contains the functional units to compute the instruction result. The functional units may be pipelined to reduce cycle time. Computed results are put on a common data bus, in the fourth stage. The common data bus is snooped on by the reservation station as was described above, and the reorder buffer, also located in the completion stage, gathers all results in a buffer. The results leave the reorder buffer in program order to be written back in the register file in the fifth stage. The fifth stage therefore is called the write back stage.

The chapter proceeds with some notes on the notation, an introduction to a half-uniary number format frequently used throughout this thesis and a detailed description of the pipeline stages.

3.1 Notation

3.1.1 Control Signals

Control signals are denoted by alphanumeric names set in a special font, like $\text{test} \in \{0, 1\}$. Control signals are not limited to bit signals, they may form multi-dimensional arrays. In this case, indices are attached to the signals according to their dimension. If we have, for example, an n -bit signal bus called data , declared in notation by $\text{data}_* \in \{0, 1\}^n$, we might reference it such:

data_i	bit i from data
$\text{data}_{j..i}$	bits j to i (in that order) from data
data_*	all the bits from data

The following example demonstrates the reference to a $m \times n$ -bit signal array / matrix called ar , declared in notation by $\text{ar}_{*,*} \in \{0, 1\}^{m \cdot n}$. This signal array is said

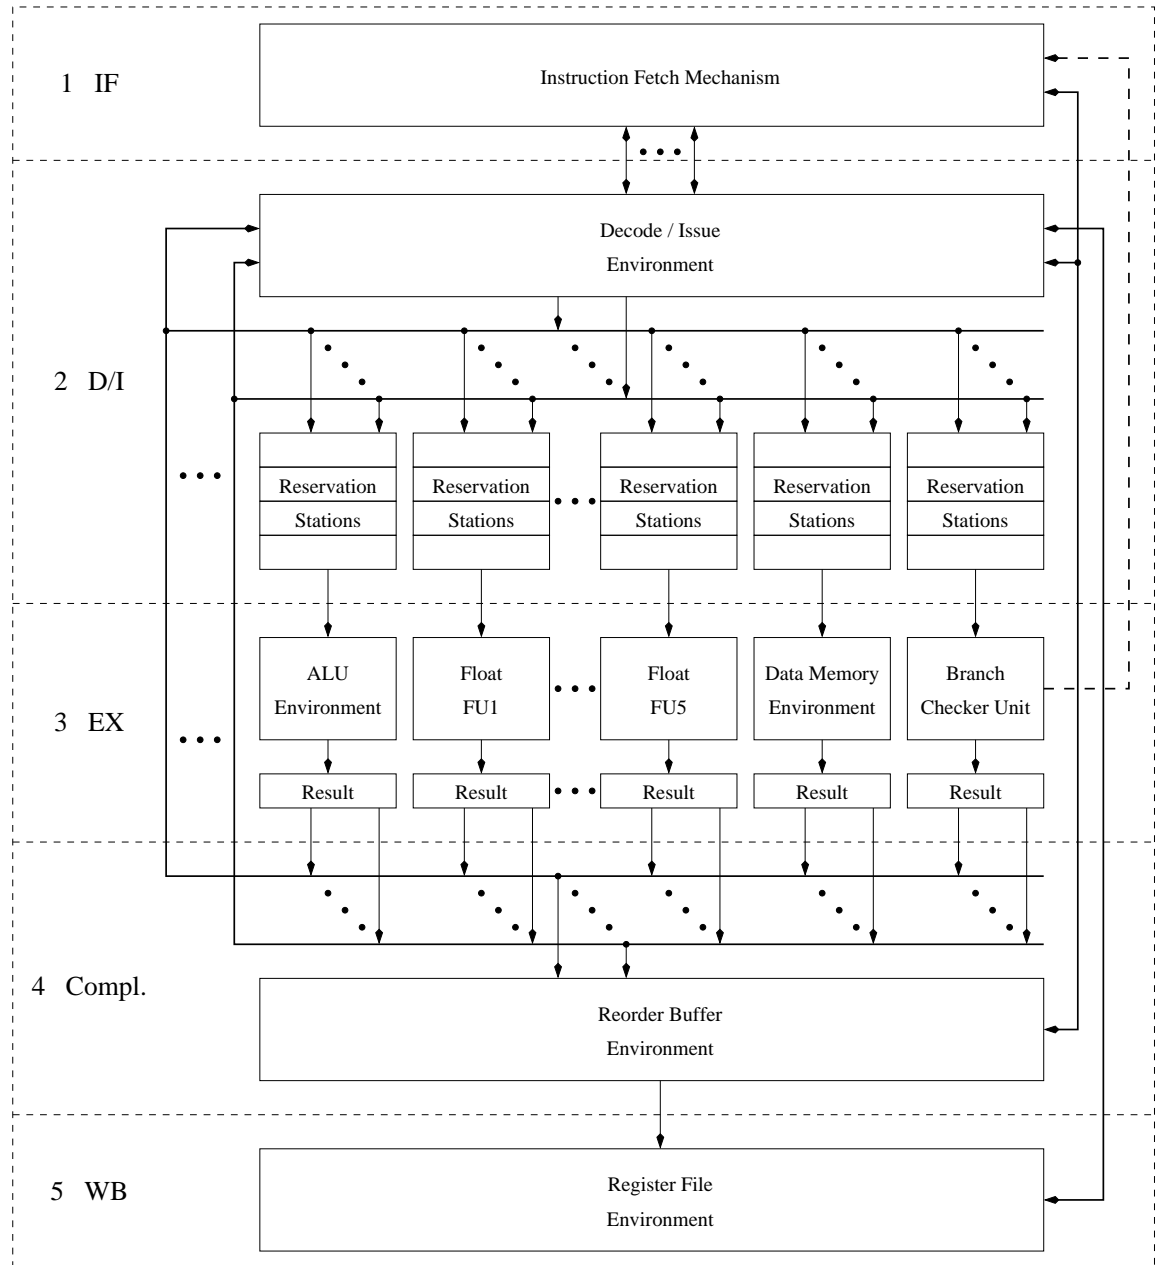


Figure 3.1: Data paths of the superscalar DLX processor

to have m rows and n columns:

$ar_{i,j}$	bit j from the i -th row of ar
$ar_{i,j..k}$	bits j to k from the i -th row of ar
$ar_{i,*}$	the i -th row of ar
$ar_{*,j}$	the j -th column of ar
$ar_{*,*}$	all the bits from ar

The star notation and the range notation is also frequently used in equations. It is a shortcut notation for using an all quantifier. The following sample notational equivalences hold for $a_*, b_*, c_*, s_* \in \{0, 1\}^n$ and $\circ \in \{\wedge, \vee\}$:

$$\begin{aligned}
 c_* = a_* &\iff \forall 0 \leq n' < n : c_{n'} = a_{n'} \\
 c_* = \overline{a_*} &\iff \forall 0 \leq n' < n : c_{n'} = \overline{a_{n'}} \\
 c_* = a_* \circ b_* &\iff \forall 0 \leq n' < n : c_{n'} = a_{n'} \circ b_{n'} \\
 c_* = (s_* ? a_* : b_*) &\iff \forall 0 \leq n' < n : (s_{n'} ? a_{n'} : b_{n'})
 \end{aligned}$$

Constructions like

$$\begin{aligned}
 c_{j..i} &= a_{(j+c)..(i+c)} \\
 c_* &= \text{test} \wedge a_* \\
 ar_{i,*} &= a_*
 \end{aligned}$$

are also allowed if the signals are properly typed; again, inserting an all quantifier reveals their meaning.

The standard order for *writing signals* is from high to low. For example, the definition $b_{3..0} := (1, a_{3..2}, 0)$ results in

$$(b_3, b_2, b_1, b_0) := (1, a_3, a_2, 0)$$

3.1.2 Busses

Often it is convenient to identify a group of signals by a single name. In our notation such signals share the same prefix ending in a dot.

An instruction operand, for instance, might consist of a valid flag $op.\text{flag} \in \{0, 1\}$, a data bus $op.\text{data}_* \in \{0, 1\}^{32}$ and a tag $op.\text{tag}_* \in \{0, 1\}^5$. These three items form a bus, abbreviated by $op.*$.

Busses themselves might also be indexed to form vectors of busses. For example, having two instruction operands denoted by op_1 and op_2 we might reference such:

$op_*.*$	both complete busses
$op_*.valid$	the valid flags of both operands
$op_*.data_*$	the data items of both operands
$op_*.tag_0$	bit 0 of the tags of both operands

If the context is clear, bus items are accessed without providing the bus prefix.

3.1.3 Figures

Figure 3.2 denotes the symbols used for gates in drawing circuits. All figures showing circuits are bounded by a dotted rectangle denoting the circuit's *interface*. Signals appearing left of or over the rectangle are *inputs* of the environment. Signals appearing right of or below the rectangle are *outputs* of the environment. Signals that




Figure 3.2: Symbols for the basic gates

have been named in a figure may be used by their name in other parts of the figure. Especially it is allowed to select single signals from an array or components of a bus by name.

The star notation and the range notation already defined will also be used in figures. For example, the signal definition $c_\star := a_\star \wedge b_\star$ we will draw as a single AND gate having inputs a_\star and b_\star and the output c_\star .

3.2 Half-Unary Number Format

This section defines the half-unary number format, which is frequently used in circuits throughout this thesis. The advantage of this format are threefold. First, the definition and proof of circuits, especially of such handling multidimensional signal arrays, is often greatly simplified by using the half-unary number format. Second, the half-unary format inherently encodes the information “not-equal-zero” and “greater-than-constant”, which often comes in handy for control definitions. Third, a half-unary encoding can be converted in constant time to a unary encoding, an encoding requiring equivalent storage space. This does not hold for the other way round. The section proceeds with a definition of the half-unary format and a description of basic properties and circuits.

Let $n \in \mathbb{N}$. The sets

$$\begin{aligned} \text{DOM}_{hu,n} &:= \{0, \dots, n\} \\ \text{BSTR}_{hu,n} &:= \{0^{n-i}1^i \mid 0 \leq i \leq n\} \end{aligned}$$

represent the domain and the valid bit strings of the n -bit half-unary encodings. Using this, we define the encoding function $\text{enc}_{hu,n}$ and the decoding (or value) function $\langle \cdot \rangle_{hu,n}$ as follows:

$$\begin{aligned} \text{enc}_{hu,n} : \text{DOM}_{hu,n} &\longrightarrow \text{BSTR}_{hu,n} \\ c &\longmapsto 0^{n-c}1^c \\ \langle \cdot \rangle_{hu,n} : \text{BSTR}_{hu,n} &\longrightarrow \text{DOM}_{hu,n} \\ 0^{n-c}1^c &\longmapsto c \end{aligned}$$

These functions are bijective and inversion functions of each other. Note that value 0 will be encoded as the zero bit string 0^n . If the index n is omitted from $\langle \cdot \rangle_{hu,n}$, it is determined by operand length.

Properties Have $\mathbf{a}_\star = 0^{n-c}1^c \in \text{BSTR}_{hu,n}$. The following properties can be easily proven:¹

$$\langle \mathbf{a}_\star \rangle_{hu} \geq i \iff \mathbf{a}_{i-1} = 1 \quad (3.1)$$

$$\langle \mathbf{a}_\star \rangle_{hu} \neq 0 \iff \mathbf{a}_0 = 1 \quad (3.2)$$

$$\langle \mathbf{a}_\star \rangle_{hu} \in \{i, i+1, \dots, j\} \iff \mathbf{a}_{i-1} = 1 \wedge \mathbf{a}_j = 0 \quad (3.3)$$

$$\langle \mathbf{a}_\star, 1^i \rangle_{hu} = \langle \mathbf{a}_\star \rangle_{hu} + i \quad (3.4)$$

$$\langle 0^i, \mathbf{a}_{n-1..i} \rangle_{hu} = \max\{0, \langle \mathbf{a}_\star \rangle_{hu} - i\} \quad (3.5)$$

$$\langle \overline{\mathbf{a}_{0..n-1}} \rangle_{hu} = n - \langle \mathbf{a}_{n-1..0} \rangle_{hu} \quad (3.6)$$

$$\langle \text{markLastOne}(\mathbf{a}_\star) \rangle_{un} = \langle \mathbf{a}_\star \rangle_{hu} \quad (3.7)$$

Remarks. The Properties 3.1 and 3.2 provide a simple means to compare half-unary encodings to constants and to 0 especially. Property 3.3 can be used to check an encoding against a given constant interval.

Properties 3.4, 3.5 and 3.6 are arithmetic properties. They justify the implementations of adders and subtractors given below.

Property 3.7 is frequently used for conversion of a half-unary encoding into a unary encoding. We define for notational convenience:

$$\text{markLastOne}(\mathbf{a}_\star) := \text{the lower } n \text{ bits of } (0, \mathbf{a}_\star) \wedge \overline{(\mathbf{a}_\star, 0)}$$

The property then is explained by the fact that the $\text{markLastOne}(\cdot)$ function simply find the *edge*, i.e. the 0-1-transition, in the half-unary encoding \mathbf{a}_\star . The functions $\text{markLastZero}(\cdot)$, $\text{markFirstZero}(\cdot)$ and $\text{markFirstOne}(\cdot)$ are defined in analogy and will be also used in this thesis.

Conversion the other way round, i.e. from a unary encoding to a half-unary encoding, can be achieved by a parallel-prefix OR circuit or, preferably for delay reasons, by a find-first-one half-unary circuit FF1hu_n described in section C.1. Note that this conversion *inherently* requires logarithmic delay. This can be shown by a reduction of the conversion function to the 1-threshold function ($\text{th1}(i_\star) = 1 \iff \exists j : i_j = 1$).² A circuit for the 1-threshold function has cost $\Omega(n)$ (a proof of this classical result can found [Weg87]) and therefore, being a 1-output function, delay $\Omega(\log n)$.

Maximum and Minimum. The maximum operation on two half-unary encodings \mathbf{a}_\star and \mathbf{b}_\star can be computed in constant time by a slice of OR gates:

$$\mathbf{c}_\star := \mathbf{a}_\star \vee \mathbf{b}_\star \implies \mathbf{c}_\star \in \text{BSTR}_{hu,n} \wedge \langle \mathbf{c}_\star \rangle_{hu} = \max\{\langle \mathbf{a}_\star \rangle_{hu}, \langle \mathbf{b}_\star \rangle_{hu}\} \quad (3.8)$$

A generalization of the maximum function on m arguments $\mathbf{a}_{0,\star}, \dots, \mathbf{a}_{m-1,\star}$ needs logarithmic time by using trees of OR gates:

$$\mathbf{c}_\star := \bigvee_{0 \leq i < m} \mathbf{a}_{i,\star} \implies \mathbf{c}_\star \in \text{BSTR}_{hu,n} \wedge \langle \mathbf{c}_\star \rangle_{hu} = \max_i \{\langle \mathbf{a}_{i,\star} \rangle_{hu}\} \quad (3.9)$$

Dually, the minimum operation on half-unary encodings can be computed using AND gates. We have:

$$\mathbf{c}_\star := \mathbf{a}_\star \wedge \mathbf{b}_\star \implies \mathbf{c}_\star \in \text{BSTR}_{hu,n} \wedge \langle \mathbf{c}_\star \rangle_{hu} = \min\{\langle \mathbf{a}_\star \rangle_{hu}, \langle \mathbf{b}_\star \rangle_{hu}\} \quad (3.10)$$

$$\mathbf{c}_\star := \bigwedge_{0 \leq i < m} \mathbf{a}_{i,\star} \implies \mathbf{c}_\star \in \text{BSTR}_{hu,n} \wedge \langle \mathbf{c}_\star \rangle_{hu} = \min_i \{\langle \mathbf{a}_{i,\star} \rangle_{hu}\} \quad (3.11)$$

¹For illustration only, drawing half-unary encodings as a white box of width $(n-c)$ followed by a black box of width c has often proved useful.

²The lowest bit of the conversion function computes the 1-threshold function.

Relations. The equality relation “=” can be computed in logarithmic time using an equality tester, because we have a unique encoding of numbers. The less-than relation “<” can be computed in logarithmic time:

$$r := \bigvee_i \overline{a_i} \wedge b_i = 1 \iff \langle a_\star \rangle_{hu} < \langle b_\star \rangle_{hu} \quad (3.12)$$

Addition and Subtraction. To compute the half-unary encoding representing the sum of the values of two other encodings, we use the following identity:

$$\langle a_\star \rangle_{hu} + \langle b_\star \rangle_{hu} = \max \{ \langle a_\star \rangle_{hu} + i \mid 0 \leq i \leq n \wedge \langle b_\star \rangle_{hu} \geq i \} \quad (3.13)$$

This equation can be directly transformed in a recipe for building a half-unary adder rewriting it with equations 3.1, 3.4, 3.9 and formally setting $b_{-1} := 1$:

$$c_\star := \bigvee_{0 \leq i \leq n} ((0^{2n-i}, a_\star, 1^i) \wedge b_{i-1}) \implies \langle c_\star \rangle_{hu} = \langle a_\star \rangle_{hu} + \langle b_\star \rangle_{hu} \quad (3.14)$$

For subtraction we use the negation property 3.6 and the adding equation 3.14. So, we have:

$$c_\star := \bigvee_{0 \leq i \leq n} ((0^{2n-i}, a_\star, 1^i) \wedge \overline{b_{n-i-1}}) \implies \langle c_\star \rangle_{hu} = n + \langle a_\star \rangle_{hu} - \langle b_\star \rangle_{hu} \quad (3.15)$$

Note that for a_\star and b_\star having different length it is advisable for delay reasons that b_\star is the shorter operand.

The Instruction Fetch Stage

3.3 Instruction Fetch Mechanism

3.3.1 Overview

This section describes the instruction fetch mechanism implemented in our processor. The fetch mechanism is a specifically tailored version of algorithm 11, p. 29, for the DLX architecture.

The fetch mechanism maintains two instruction buffers, called instruction fetch queues (IFQs). In each cycle one of these queues is designated the *fetching IFQ* and one, not necessarily different, queue is designated the *issuing IFQ*. The fetching IFQ receives newly fetched instructions and from the issuing IFQ fetched instruction are being issued.

Initially the fetching IFQ and the issuing IFQ are equal. The situation changes, when the control predicts or resolves a control flow change from a CFI in the issuing IFQ. Then, the fetching IFQ will be switched and fetched instruction will be stored in the alternative queue. The issuing IFQ follows this switch, if the instruction causing the control flow change has been issued.

On control flow speculation, the fetch mechanism goes in the outstanding speculation state. It will not predict or resolve further CFIs, until the *branch checker unit* (BCU) has positively or negatively verified the prediction. The report of the BCU causes the clearance of the outstanding speculation state. On a misprediction, the fetch mechanism has to roll back to the corrected target, involving initialization of the fetching and issuing IFQ and setting its fetch PC anew. For the special case, that a predicted taken branch is corrected to a fall-through branch, the alternative IFQ still contains the instructions following a branch; i.e. only a switch of IFQs without initialization is required.

```

1: if more than  $k$  free entries in fetchIFQ then
2:   Request  $k$  instructions from memory starting at fetchIFQ.PC
3:   Push the  $k$  instructions from the memory on fetchIFQ
4:   fetchIFQ.PC  $\leftarrow$  fetchIFQ.PC +  $k$ 
5: end if

```

Algorithm 12: Implementation algorithm for instruction fetch

```

1: if rollback then
2:   if JISR then
3:     Switch fetchIFQ
4:     Initialize fetchIFQ
5:     Set issueIFQ to fetchIFQ
6:     fetchIFQ.PC  $\leftarrow$  SISR
7:   else
8:     IWdisabled  $\leftarrow$  0
9:   end if
10: end if

```

Algorithm 13: Implementation algorithm for rollback

Exceptions may also interrupt regular instruction fetch. Similar to the treatment of mispredictions, fetching and issuing IFQ are switched and cleared and the fetch PC is set to the start of the interrupt service routine.

Buffering instructions in prefetch queues is a common approach in modern microprocessors ([Bha96, WS94]). The alternating prefetch queue design can also be found in the Pentium processor, although [AS95] does not provide an in-depth discussion.

3.3.2 Algorithms

The details of the instruction fetch mechanism are given in five algorithms dealing with the fetch of instructions, the misprediction treatment, the interrupt treatment, the prediction of control flow and the construction of the instruction window. Each round, the algorithms are executed in that sequence.

Algorithm 12 shows the implementation algorithm for instruction fetch. Instruction fetch is made in blocks of k subsequent instructions. Therefore it stalls, if the fetching IFQ has no room for k instructions left (l. 1). The instruction requested and fetched are pushed in the fetching IFQ (ll. 2–3). After this, the fetch PC advances by k steps (l. 4).

Algorithm 13 shows the implementation algorithm for the treatment of rollback. Rollback is initiated by the global signal rollback (l. 1), generated during retirement. In case that JISR is activated in addition (l. 2), the fetch IFQ is switched and initialized (ll. 3–4), followed by the issuing IFQ (l. 5). The interrupt service routine is called (l. 6). If rollback is activated but not JISR (l. 7), a mispredicted CFI was caught retiring. Since in this case the branch checker unit will have reported a misprediction before, the control flow has already been changed adequately by the misprediction implementation algorithm described below. Only the instruction window, temporarily disabled for issuing, has to be enabled again by zeroing the status variable IWdisabled (l. 8).

Algorithm 14 on the next page shows the implementation algorithm for the treatment of reports from the branch checker unit. The branch checker unit verifies predictions made by the fetch mechanism. Its implementation is described in detail in section 3.11.4. For the moment, it is sufficient to present the branch checker unit

Component	Width	Purpose
valid	1	indicates valid bus contents
mp	1	indicates misprediction
jump	1	indicates jump instruction
btaken	1	indicates taken branch
cfcPC _*	32	corrected CFC address

Table 3.1: Branch checker unit bus

```

1: if bcu.valid then
2:   oss  $\leftarrow$  0
3:   if bcu.mp then
4:     Switch fetchIFQ
5:     Set issueIFQ to fetchIFQ
6:     if bcu.jump  $\vee$  bcu.btaken then
7:       Initialize fetchIFQ
8:       fetchIFQ.PC  $\leftarrow$  bcu.cfcPC
9:     end if
10:    IWdisabled  $\leftarrow$  1
11:  end if
12: end if

```

Algorithm 14: Implementation algorithm for BCU treatment

interface. The BCU reports its verifications on a bus `bcu.*`. Table 3.1 lists the components. Signal `bcu.valid` indicates that a verification took place. On a misprediction, `bcu.mp = 1`, two cases are distinguished: for a jump instruction, the corrected PC `cfcPC*`; for a branch instruction, `btaken` indicates the correct outcome, `cfcPC*` will also contain the correct PC.

We continue the description of algorithm 14. The algorithm first checks if the BCU reports a valid bus (l. 1). Any report, be good or ill, results in clearance of the outstanding speculation state (l. 2). A misprediction of control flow results in the initiation of a rollback (ll. 3–11). The fetch IFQ is switched (l. 4), the issuing IFQ follows (l. 5). If the mispredicted CFI was a branch and has been falsely taken, `fetchIFQ` still holds the instruction of the fall-through target and need not be initialized. Otherwise, the (new) fetch IFQ is initialized (l. 7) and its PC is set to the corrected address, as reported by the BCU (l. 8). Note that branch checking is done out of order to speed up the machine: fetch of the correct control flow may start as soon as the misprediction is detected, though the execution of the correct control flow start after machine rollback. Mispredicted CFIs are therefore passed on to retirement and will generate a regular rollback call there if the control flow has not been interrupted by a previous interrupt condition. Instruction issue is disabled by setting `IWdisabled` until then (l. 10). The rollback algorithm will enable the instruction window again, as has already been seen above.

Algorithm 15 shows the implementation algorithm used for prediction and resolving of CFIs. The issuing IFQ is scanned for CFIs in case that the machine is not in an outstanding speculation situation and its instruction window is not disabled (l. 1). The first found CFI is renamed to `cfi` for ease of notation (l. 2). If `cfi` can be resolved or predicted, the algorithm will do so in ll. 4–17. For a branch, the intermediary variable `cfi.cfcPC` is set to the relative jump target (l. 5) and the variable `cfi.cfc`, indicating a control flow change, is set to the predicted or resolved outcome. Otherwise, the CFI is a jump instruction. In this case, `cfi.cfcPC` is set to the predicted or resolved target (l. 8) and `cfi.cfc` is set (l. 9) since the DLX jump instruction

```

1: if  $\overline{oss} \wedge \overline{IWdisabled} \wedge \exists j : \text{issuelFQ}_j \text{ is CFI}$  then
2:   Let  $\text{cfi}$  denote  $\text{issuelFQ}_j$ ,  $j$  minimal s.t.  $\text{issuelFQ}_j$  is CFI
3:   if can predict / resolve  $\text{cfi}$  then
4:     if  $\text{cfi}$  is branch then
5:        $\text{cfi.cfcPC} \leftarrow \text{cfi.PC} + \text{cfi.imm16}$ 
6:        $\text{cfi.cfc} \leftarrow (\text{predicted} / \text{resolved taken} ? 1 : 0)$ 
7:     else
8:        $\text{cfi.cfcPC} \leftarrow \text{predicted} / \text{resolved target}$ 
9:        $\text{cfi.cfc} \leftarrow 1$ 
10:    end if
11:    On prediction only:  $\text{oss} \leftarrow 1$ 
12:    if  $\text{cfi.cfc}$  then
13:      Switch  $\text{fetchIFQ}$  and initialize
14:       $\text{fetchIFQ.PC} \leftarrow \text{cfi.cfcPC}$ 
15:    end if
16:     $\text{cfi4l.valid} \leftarrow 1$ 
17:     $\text{cfi4l.cfc} \leftarrow \text{cfi.cfc}$ 
18:     $\text{cfiReady} \leftarrow 1$ 
19:  end if
20: end if

```

Algorithm 15: Implementation algorithm for prediction and resolving

```

1: if  $\overline{IWdisabled}$  then
2:   Set  $\text{issuelFQ}_j.\text{valid} \leftarrow 1$  for all  $j$ 
3:   if  $\exists j : \text{issuelFQ}_j \text{ is CFI}$  then
4:     Let  $j$  be minimal s.t.  $\text{issuelFQ}_j$  is CFI
5:     Set  $\text{issuelFQ}_{j'}.\text{bbi} \leftarrow 1, \text{issuelFQ}_{j'}.\text{cfi} \leftarrow 0$  for all  $j' < j$ 
6:     Set  $\text{issuelFQ}_j.\text{bbi} \leftarrow 0, \text{issuelFQ}_j.\text{cfi} \leftarrow \text{cfi4l.valid}$ 
7:     Set  $\text{issuelFQ}_{j'}.\text{bbi} \leftarrow 0, \text{issuelFQ}_{j'}.\text{cfi} \leftarrow 0$  for all  $j' > j$ 
8:   else
9:     Set  $\text{issuelFQ}_j.\text{bbi} \leftarrow 1$  for all  $j$ 
10:  end if
11: else
12:   Set  $\text{issuelFQ}_j.\text{valid} \leftarrow 0$  for all  $j$ 
13: end if
14: Wait for issue
15: Drain the issued instructions from the issuing IFQ.
16: if  $\exists j : \text{issuelFQ}_j \text{ issued} \wedge \text{issuedIFQ}[j].\text{cfi}$  then
17:    $\text{cfi4l.valid} \leftarrow 0$ 
18:   if  $\text{cfi4l.cfc}$  then
19:     Switch the issuing IFQ
20:   end if
21: end if

```

Algorithm 16: Implementation algorithm for instruction window construction

always cause a control flow change. For a prediction, the fetch mechanism sets the outstanding speculation state, preventing further predictions until the BCU signals verification (l. 11). In case of control flow change, predicted or not, the fetching IFQ is switched and its PC is set to the cfi.cfcPC variable (ll. 13–14). Finally, setting cfi4l.valid indicates that a CFI awaits issuing (l. 16). Variable cfi4l.cfc indicates that the control flow breaks at this instruction (l. 17).

Name	Width	Description
IFQ _i .PC _*	32	fetch PC for IFQ _i , $i \in \{0, 1\}$
cfcPC _*	32	control flow change PC buffer
fetchIFQ.idx	1	index of the fetching IFQ
issueIFQ.idx	1	index of the issuing IFQ
oss	1	outstanding speculation
IWdisabled	1	disabled instruction window
cfi4l.valid	1	indicate that CFI is ready to be issued
cfi4l.cfc	1	indicates that next CFI causes a CFC
P _*	7	pending action

Table 3.2: Registers in the instruction fetch stage

Algorithm 16 presents the procedure to construct the instruction window for the Tomasulo machine core. In case that the instruction window is not disabled (l. 1), the instructions in the issuing IFQ are set valid. First assume that a CFI is present in the issuing IFQ (l. 3–7). Basicblock instructions before the first CFI are marked as such (l. 5). The CFI will be marked as non-bbi; it is marked as CFI iff it is ready for issue, `cfi4l.valid` = 1 (l. 6). The following instructions are marked neither BBI or CFI, so that the Tomasulo core will not issue them (l. 7). If no CFI exists, all the instructions stored in the queue are basicblock instruction and marked appropriately (l. 9). A disabled instruction window will be all set to invalid instructions (l. 12).

After this preparation, instruction issue can take place (l. 14). The issued instructions will be drained from the issuing IFQ (l. 15). If these instructions included a CFI (l. 16), the `cfi4l.valid` variable is cleared (l. 17), and the issuing IFQ is switched if the cfi involved a control flow change (ll. 18–20).

3.3.3 Control

In hardware, the implementation algorithms are parallelized. The execution of the algorithms is subject to the conventions of the memory protocol. As this protocol forbids the interruption of read requests, the control flow can only be changed after completion of the previous read request. So, the control must reflect that the processor cycles are superimposed by the “memory cycles” (as indicated by the `IM.busy` signal).

Registers

We proceed with the description of the registers that the instruction fetch stage is maintaining. Most of these register appeared already as variables in the algorithms. Table 3.2 shows an overview; details follow:

- Registers IFQ₀.PC_{*} and IFQ₁.PC_{*} store a fetch PC of IFQ₀ and IFQ₁ respectively. Register `cfcPC*` is the scrapbook register for buffering the PC of control flow changes.

Update of these three register is performed in the PC environment, described in section 3.6.

- Registers `fetchIFQ.idx` and `issueIFQ.idx` store the index of the fetching and of the issuing IFQ respectively.

Updates of these registers are controlled by three signals. Signal `flfqswitch` and `ilfqswitch` indicates the toggling of `fetchIFQ.idx` and `issueIFQ.idx`. Signal

#	Name	Detection	Active Control Signals	
			on Enter	on Leave
0	CFC	$\text{cfi.ready} \wedge \text{cfi.cfc}$	cfcPC.ce	flFQswitch
1	MP0	$\text{bcu.valid} \wedge \text{bcu.mp}$	cfcPC.ce disableIW	$\text{flFQswitch}, \text{ilFQ2flFQ}$
2	MP0 _{sw}	$\text{bcu.valid} \wedge \text{bcu.mp}$ $\wedge (\text{bcu.jump} \vee \overline{\text{bcu.btaken}})$	cfcPC.ce disableIW	$\text{flFQswitch}, \text{ilFQ2flFQ}$ dontInitlFQ
3	MP1	rollback	enableIW	
4	MP2	$\text{MP0} \wedge \text{rollback}$	—	$\text{flFQswitch}, \text{ilFQ2flFQ}$ enableIW
5	MP2 _{sw}	$\text{MP0sw} \wedge \text{rollback}$	—	$\text{flFQswitch}, \text{ilFQ2flFQ}$ $\text{dontInitlFQ}, \text{enableIW}$
6	JISR	JISR	cfcPC.ce disableIW	$\text{flFQswitch}, \text{ilFQ2flFQ}$ enableIW

Table 3.3: Pending actions in the instruction fetch mechanism

ilFQ2flFQ indicates that the issuing IFQ index will be set to the fetching IFQ index. We have:

$$\begin{aligned} \text{fetchIFQ.idx}' &:= \text{fetchIFQ.idx} \oplus \text{flFQswitch} \\ \text{issueIFQ.idx}' &:= (\text{ilFQ2flFQ} ? \text{fetchIFQ.idx}' : \text{issueIFQ.idx} \oplus \text{ilFQswitch}) \end{aligned}$$

- Register oss indicates if the fetch mechanism is in an outstanding speculation situation.

As the algorithms state, the oss register is set on a prediction of a CFI and cleared on rollback, misprediction and successful verification. We have:

$$\text{oss}' := (\text{oss} \wedge \overline{\text{bcu.valid}}) \vee (\text{cfi.ready} \wedge \text{cfi.doPred})$$

- Register IWdisabled indicates a disabled instruction window. It is updated by the control signals disableIW and enableIW according to the following equation:

$$\text{IWdisabled}' := (\text{IWdisabled} \vee \text{disableIW}) \wedge \overline{\text{enableIW}}$$

- Register cfi4l.valid is active, if a CFI in the issuing IFQ awaits issue. It will be set, if a CFI is ready and it will be reset, if a CFI is issued or the instruction window has been disabled:

$$\text{cfi4l.valid}' := (\text{cfi4l.valid} \vee \text{cfi.ready}) \wedge \overline{(\text{cfilssued} \vee \text{disableIW})}$$

If $\text{cfi4l.valid} = 1$ register cfi4l.cfc denotes, if the CFI awaiting issue generates a change control flow. The register is updated with cfi.cfc , whenever the prediction environment got a prediction ready, signalled by cfi.cfc :

$$\text{cfi4l.cfc}' := (\text{cfi.ready} ? \text{cfi.cfc} : \text{cfi4l.cfc})$$

- The pending action register $P_\star \in \{0, 1\}^7$ will be described in detail below.

Pending Actions

Changes to the fetching IFQ may be only allowed on the completion of an instruction memory read access. There are three actions of the algorithms resulting in a change of the fetching IFQ (and the fetch PC):

- The prediction or resolving of a CFI causing a control flow change (cf. algorithm 15, ll. 12–15).
- A misprediction report by the branch checker unit (cf. algorithm 14, ll. 4–9). The instruction fetch mechanism may *initiate* rollback but might not complete until a rollback is signalled from the retirement stage as well.
- A JISR condition is reported by the retirement stage (cf. algorithm 13, ll. 3–6). A rollback to the interrupt service must be initiated.

Since all the three actions might possibly be detected in a single memory cycle, they must be prioritized to ensure the same order of execution as in the algorithms. Regular control flow changes are overruled by mispredictions and JISR conditions, misprediction is overruled by JISR conditions. Furthermore, the misprediction case has to be broken down in three subcases, to treat the waiting for the rollback adequately.

Table 3.3 shows an overview of the events we catch in each memory round. They are ordered by *increasing* priority. Each event has a detection signal associated, which evaluates to 1 if the preconditions for the event are met. An event will “enter” its pending execution, if it is the event of highest priority detected in the current memory round. On entering control signals as listed will be activated. If the instruction memory is not busy, as indicated by $\overline{\text{IM}}.\text{busy}$, the event of highest priority so far detected may actually be executed; it “leaves” the state of pending action and activates the signals listed under column “on leave”. We continue with a detailed description of the cases and will then develop a circuit to control the pending actions.

Following are the descriptions of the cases listed in table 3.3 from low to high priority:

- The prediction environment (cf. section 3.7), scanning for CFIs, signals $\text{cfi.ready} \wedge \text{cfi.cfc}$ if a control flow instruction causing a control flow change has been resolved or predicted.

On enter, the determined cfcPC_* is clocked by setting cfcPC.ce . On leave, the fetching IFQ is switched by setting fifQswitch .

- Misprediction treatment must be broken down into a number subcases according to the two phases detection and retirement of a mispredicted CFI.

The cases MP0 and MP0sw account for the detection of the misprediction by the report $\text{bcu.valid} \wedge \text{bcu.mp} = 1$ of the branch checker unit. If additionally $\text{bcu.jump} \vee \text{bcu.btaken} = 1$, the rollback will be by switching the IFQ only. On enter, MP0 and MP0sw both clock the cfcPC_* register. On leave, both will switch the fetching IFQ and set the issuing IFQ to the fetching IFQ. MP0sw will prevent the initialization of the fetching IFQ by setting dontInitIFQ . Additionally, both will disable the instruction window by the signal disableIW .

Cases MP1, MP2 and MP2sw deal with the treatment of the rollback, if the mispredicted CFI is retired in the reorder buffer. At earliest, this is one cycle after the detection of cases MP0 or MP0sw.

Case MP2 and MP2sw deal with the rollback of a mispredicted CFI in the memory cycle *of the detection of MP0 or MP0sw*. Their detection therefore depends on the detection of MP0 or MP0sw in the same cycle. They will perform the same updates as MP0 and MP0sw and additionally enable the instruction window (enableIW) on leave.

Case MP1 deals with the mispredicted CFI retiring in a memory cycle *after detection of MP0 or MP0sw*. The queue states therefore has already been




Figure 3.3: Pending actions of the instruction fetch

updated appropriately and only the instruction window needs to be enabled again.

- Detection of an interrupt requests by $JISR = 1$ has the highest priority. It disables the instruction window on entering by activating $disableIW$ and clocks the start of the interrupt service routine in the $cfcPC_*$ register by setting $cfcPC.ce$. On leave, the fetching IFQ is switched by setting $fIFQswitch$ and the issuing IFQ follows ($iIFQ2fIFQ$). The instruction window is enabled again.

Figure 3.3 shows the circuit used to implement this control. The register P_* holds the current action level detected in half-unary encoding.

Arriving at the control are the detection signals $P_*.detect$ from table 3.3.³ With a half-unary maximum computation we filter out the detected action:⁴

$$detectMax_* := \bigvee_{0 \leq i < 7} (0^{7-i}, P_i.detect^i)$$

The new (auxiliary) pending level can likewise be computed by a half-unary maximum computation:

$$auxP_* := P_* \vee detectMax_*$$

Now, the $auxP_*$ is also the new value of the P_* register if the memory is still busy fetching. Otherwise the pending register will be cleared:

$$P'_* := auxP_* \wedge IM.busy$$

The unary signals $enter_*$ and $leave_*$ signal the entering and leaving of a certain pending level. We have $enter_i = 1$ iff level i was entered in the round and $leave_i = 1$ iff level i was left in the round. We define as follows:

$$\begin{aligned} enter_* &:= markLastOne(auxP_*) \wedge \overline{P_*} \\ leave_* &:= markLastOne(auxP_*) \wedge \overline{IM.busy} \end{aligned}$$

³Define $MP0 := P_1 \wedge \overline{P_2}$ and $MP0sw := P_2 \wedge \overline{P_3}$

⁴Actually this computation simplifies greatly on filling in the detection expressions. This strategy will not be followed here for ease of reading.

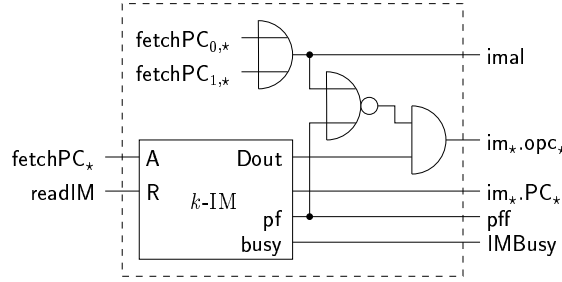


Figure 3.4: Instruction memory environment

These equations are explained as follows: the highest pending level, given in half-ary encoding by $\text{markLastOne}(\text{auxP}_*)$ is entered, if it is greater than the previous pending level. It is left, if the instruction memory is not busy.

The control signals given in table 3.3 can be computed via the enter_* and leave_* busses. If S is such a signal, we set

$$S := \bigvee_{\substack{\text{s activated on enter } i}} \text{enter}_i \vee \bigvee_{\substack{\text{s activated on leave } i}} \text{leave}_i$$

3.4 Instruction Memory Environment

The design of the instruction memory environment IMenv follows the design presented in [Krö99]. However, the underlying instruction memory $k\text{-IM}$ always returns k instruction opcodes and their PCs starting at the base address $\langle A_* \rangle_2$, not necessarily a multiple of k . The design of such instruction memory is not treated here; such extensions are closely related to cache design. An approach to cache design can be found in [MP00].

Figure 3.4 shows the implementation of the IMenv . Read requests are controlled by the IM read signal readIM . This signal requests reading whenever there is enough room in the fetching IFQ:

$$\text{readIM} := \overline{(\text{fetchIFQ.idx} ? \text{ifq}_1.\text{full} : \text{ifq}_0.\text{full})}$$

The fetch PC fetchPC_* , forwarded from the PC environment, specifies the read address. On IM.busy Instruction PCs are returned on the busses $\text{im}_*.PC_*$. Instruction opcodes are returned on the busses $\text{im}_*.opc_*$. In two cases the opcodes are cleared, resulting in a nop instruction having no effect: First, the instruction fetch may be misaligned, signalled by $\text{im}_*.imal := \text{fetchPC}_1 \vee \text{fetchPC}_0$. Second, one of the instructions may reside on a memory page swapped out of physical memory. In that case, the memory returns a page fault signal. This convention requires an additional software requirement for the interrupt service routine: on return from page fault recovery it must be guaranteed that all k instructions starting at address $\langle \text{fetchPC}_* \rangle_2$ reside in physical memory. The page fault signal is returned on $\text{im}_*.pff$.

3.5 Instruction Fetch Queue Environment

Figure 3.5 shows the implementation of the instruction fetch queue environments using a multiported queue. Multiported queues are data structures that allow storing and retrieving of multiple elements in a round on a first-in-first-out basis. See

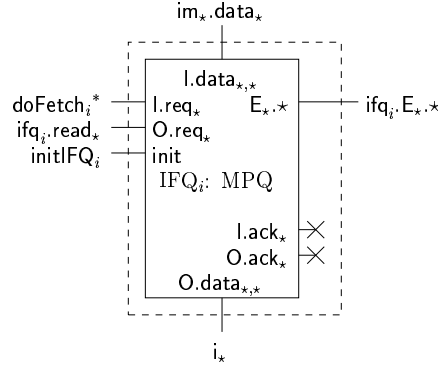


Figure 3.5: Instruction fetch queue environment

Component	Width	Purpose
valid	1	indicates valid entry (maintained by queue control)
PC _*	32	instruction's program counter
opc _*	32	instruction's opcode
pff	1	indicates page fault on fetch
imal	1	indicates misalignment

Table 3.4: Components of the instruction fetch queue entries

section 5.3 for the abstract definition, implementation and correctness proof of a multiported queue.

Table 3.4 shows the composition of an entry in the IFQ. The PC_{*} and opc_{*} component are self-explanatory. The page fault exception bit pff and the misalignment exception bit imal are passed by the instruction memory environment.

Each IFQ is controlled by three signals. The signal initIFQ_i requests an initialization (i.e. a clearing of the valid bits) of the queue. The signal doFetch is activated in the same cycle that instructions from the memory arrive at this queue.

$$\begin{aligned}
 \text{doFetch}_i &:= \overline{\text{IM.busy}} \wedge \text{readIM} \wedge (\text{fetchIFQ.idx} = i) \\
 \text{initIFQ}_i &:= \text{fIFQswitch} \wedge \overline{\text{dontInitIFQ}} \wedge (\text{fetchIFQ.idx}' = i)
 \end{aligned}$$

Data is requested for read out by the read signals ifq_i.r_{*} generated by the issue selector environment. It leaves the queue on the output busses O_{*}. For prediction purposes, the queue entries are exported on the entry bus E_{*}._{*} and received back on E'_{*}._{*}.

3.6 PC environment

Figure 3.6 shows the PC environment. It contains two registers, IFQ₀.PC_{*} and IFQ₁.PC_{*}, holding the fetch PC for IFQ₀ and IFQ₁. Only one of these PCs is active at a time—in the sense of addressing the instruction memory. This PC is selected by the index of the fetching IFQ, fetchIFQ.idx, generated by the control. The resulting bus fetchPC_{*} was already used to address the instruction memory in the IMenv.

The update of these registers is as follows. For $i = \text{fetchIFQ.idx}$ register IFQ_i.PC_{*} receives the incremented-by- k -version of fetchPC_{*}. It will only store it, if a fetch takes place, as indicated by the signal doFetch_i, which has been defined already. For $i = \overline{\text{fetchIFQ.idx}}$ register IFQ_i.PC_{*} is updated if it becomes the fetching PC in

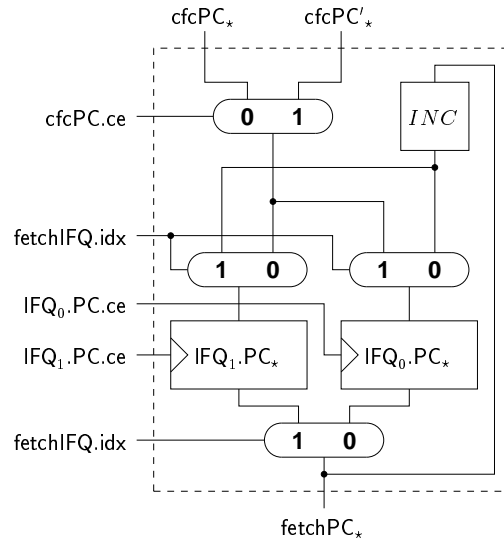


Figure 3.6: PC environment




Figure 3.7: Control flow change PC generation

Signal	opc _{31..26}	Associated instructions
t.epc	111111	rfe
t.gpr	01011*	jr, jalr
t.imm16	0001**	beqz, bnez, fbeqz, fbnez
t.imm26	00001*	j, jal
cfi	t.epc \vee t.gpr \vee t.imm16 \vee t.imm26	
bbi	cfi	
relJump	t.imm16 \vee t.imm26	
absJump	relJump	

Table 3.5: Control signals in the control flow environment

the next round and is initialized, i.e. if $\text{initIFQ}_i = 1$. In this case it is fed either the stored value cfcPC_\star or the update value cfcPC'_\star of the CFC PC register, chosen by cfcPC.ce .

In summary, the clock enable signals of these register is defined as follows:

$$\text{IFQ}_i.\text{PC.ce} := \text{doFetch}_i \vee \text{initIFQ}_i$$

3.7 Prediction Environment

The prediction environment PredEnv is part of the hardware implementation of the prediction algorithm 15. The prediction environment scans for CFI in the issuing IFQ and attempts to predict or resolve them.

The following sections review the scan procedure for the CFI, the prediction of CFIs, the resolving of CFIs and the construction of the following bus:

cfi.ready	indicate that a CFI is ready
cfi.doPred	the CFI was predicted
cfi.btaken	the CFI is a taken branch
cfi.cfc	the CFI causes a control flow change
cfi.cfcPC _★	the control flow change PC of the CFI

3.7.1 Finding the first CFI

We start with the selection of the issuing IFQ from IFQ_0 and IFQ_1 :

$$\text{issueIFQ}_\star.★ := (\text{issueIFQ.idx} ? \text{IFQ}_1.\text{E}_\star.★ : \text{IFQ}_0.\text{E}_\star.★)$$

In hardware this is done in three steps.

First, the instruction opcodes $\text{issueIFQ}_\star.\text{opc}_\star$ are predecoded to determine their type. Table 3.5 shows the signals generated for each entry in the fetching IFQ. Refer to appendix A for an overview of the DLX instruction set architecture. The signals t.epc , t.gpr , t.imm16 and t.imm26 characterize the target a CFI can take; relJump indicates PC-relative jumps as opposed to absolute jumps. The signals cfi and bbi classify instructions into basic-block and control flow instructions.

Second, a find-first-one half-unary circuit FF1hu_n receiving $\text{issueIFQ}_\star.\text{cfi}$ returns the position of the first control flow instruction in the issuing IFQ in half-unary encoding. If $\text{cfiNegHU}_\star \in \{0, 1\}^n$ is the output of the circuit, we have

$$\langle \text{cfiNegHU}_\star \rangle_{hu} = l \iff \begin{aligned} & \text{issueIFQ}_{l'}.\text{cfi} = 0 \quad l' < l \\ & \wedge \quad \text{issueIFQ}_l.\text{cfi} = 1 \end{aligned}$$

Third, the found instruction is driven on an auxiliary bus $\text{cfi}.\star$ by using

$$\text{firstCFI}_\star := \text{markFirstOne}(\text{cfiNegHU}_\star)$$

as an output enable signal. The contents of $\text{cfi}.\star$ are valid, if a candidate was found, the fetch mechanism has no outstanding speculation and no action is pending in the P_\star register:

$$\text{cfi.valid} := \text{cfiNegHU}_{n-1} \wedge (\overline{\text{oss}} \vee \text{bcu.valid} \wedge \overline{\text{bcu.mp}}) \wedge \overline{(\bigvee \text{P}_\star)}$$

3.7.2 Prediction

The $\text{cfi}.\star$ bus is passed on to the branch predictor unit (BPU) for prediction. The BPU acknowledges the prediction requests with the signal bpu.pred . On acknowledgement, the BPU guarantees a valid bpu.btaken signal indicating the outcome of a branch or a valid bpu.cfcPC_\star bus indicating an absolute jump target.

Appendix B introduces a sample branch predictor for branch instructions only (so $\text{bpu.pred} := \text{cfi.t.imm16}$).

3.7.3 Resolving

To resolve branches, we access the source operand data generation of the decode / issue environment. The bus firstCFI_\star is used to select the first source operand of the CFI by drivers from all the first operands $\text{i}_\star.\text{op}_1.\text{data}_\star$ in the decode / issue environment:

$$\text{cfi.op}_1.\text{data}_\star := \text{i}_i.\text{op}_1.\text{data}_\star \quad \text{with } \text{firstCFI}_i = 1, i \in \{0, \dots, \beta - 1\}$$

The contents of this bus are only valid, if the decode / issue environment actually “sees” the CFI, i.e. if the CFI has an index which is less than the issue width. Additionally, the data on $\text{i}_i.\text{op}_1.\text{data}_\star$ is only valid if $\text{i}_i.\text{op}_1.\text{valid} = 1$. We obtain with selecting on $\text{i}_\star.\text{op}_1.\text{valid}$ by drivers:

$$\text{cfi.op}_1.\text{valid} := \overline{\text{cfiNegHU}_\beta} \wedge \text{i}_i.\text{op}_1.\text{valid} \quad \text{with } \text{firstCFI}_i = 1, i \in \{0, \dots, \beta - 1\}$$

We resolve a CFI in three situations:

- A `rfe` instruction is resolved, when the ROB is emptied. This ensures that the exceptional registers ESR_\star and EPC_\star have the values of the last instruction writing to it.
- The unconditional relative jump instructions `j` and `jal` are resolved right-away, since their target address can be immediately computed and they are known to jump unconditionally.
- The register jump instructions `jr` and `jalr` are resolved when their first operand becomes valid. The operand contains the target address.

This conditions are summarized in the following equation:

$$\text{cfi.resolve} := (\text{cfi.t.epc} ? \text{ROB.empty} : \text{cfi.op}_1.\text{valid} \vee \text{cfi.t.imm26})$$

If $\text{cfi.resolve} = 1$, a taken branch is signalled by

$$\text{cfi.resolve.btaken} := \text{cfi.op}_1.\text{eqz} \oplus \text{cfi.opc}_{26}$$

where $\text{cfi.op}_1.\text{eqz} := (\bigvee \text{cfi.op}_1.\text{data}_\star)$. The target of a direct jump can be directly taken from $\text{cfi.op}_1.\text{data}_\star$.

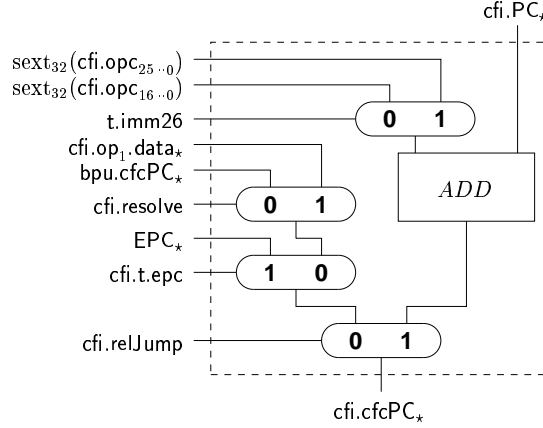


Figure 3.8: Generation of the CFI CFC PC

3.7.4 Construction of the CFI bus

A CFI is ready for execution, if it is valid can be either predicted or resolved:

$$\text{cfi.ready} := \text{cfi.valid} \wedge (\text{cfi.resolve} \vee \text{bpu.pred})$$

Resolving a CFI has priority over prediction. This is reflected in the definitions of the signals `doPred`, telling the BPU to predict actually, `cfi.btaken` indicating a taken branch and `cfi.cfc` indicating a control flow change:

$$\begin{aligned} \text{cfi.doPred} &:= \overline{\text{cfi.resolve}} \wedge \text{bpu.pred} \\ \text{cfi.btaken} &:= (\text{cfi.resolve} ? \text{cfi.resolve.btaken} : \text{bpu.btaken}) \\ \text{cfi.cfc} &:= \overline{(\text{cfi.t.imm16} \wedge \text{cfi.btaken})} \end{aligned}$$

Finally, prediction or resolving results in the computation of a control flow change PC for the CFI, `cfi.cfcPC*`. Figure 3.8 shows how to compute this bus. For relative jumps, indicated by `cfi.relJump`, the `cfcPC*` is computed as the PC of the CFI plus the immediate constant, either 16 or 26 bits. Otherwise, the target is the `EPC*` register for `rfe` instructions, the resolved or the predicted address for jumps to register locations.

3.8 Instruction Window Environment

3.8.1 Construction of the Instruction Window

The construction of the instruction window takes place in the instruction window environment `IWenv`. Here, the instructions that the fetch mechanism choses to present to the decode / issue environment are selected and the entries are padded with the necessary interface information (classification of instructions by flags `bbi`, `cfi` and providing the next PC `nPC*`) to the decode / issue environment.

Some work has already been done in section 3.7.1: recall the definition of the issuing IFQ as

$$\text{issueIFQ}_{*,*} := (\text{issueIFQ.idx} ? \text{IFQ}_1.E_{*,*} : \text{IFQ}_0.E_{*,*})$$

In accordance to the nomenclature used in the algorithms we rename the first entries of the issuing IFQ. For clarity, we explicitly list all the renamed bus items, $i \in$

$\{0, \dots, \beta - 1\}$:

$$\begin{aligned} i_i.opc_\star &:= \text{issueIFQ}_i.opc_\star \\ i_i.PC_\star &:= \text{issueIFQ}_i.PC_\star \\ i_i.iml &:= \text{issueIFQ}_i.iml \\ i_i.pff &:= \text{issueIFQ}_i.pff \end{aligned}$$

The valid, bbi, cfi signals are computed as follows:

$$\begin{aligned} i_\star.valid &:= \text{issueIFQ}_\star.valid \wedge \overline{IWdisabled} \\ i_\star.bbi &:= \overline{cfiNegHU_\star} \\ i_\star.cfi &:= \text{firstCFI}_\star \wedge (cfi.ready \vee cfi4l.valid) \\ i_\star.cfc &:= \text{firstCFI}_\star \wedge (cfi.ready ? cfi.cfc : cfi4l.cfc) \end{aligned}$$

The abstract interface of a speculative instruction fetch mechanism requires to compute the next PC for every instruction. As the PC is not a general-purpose source register in the DLX instruction set, we need to supply its alleged new value only for the verification of control flow instructions. Since at most one control flow instruction is presented to the decode / issue stage in the issuing IFQ, we supply the next PC via

$$cfi4l.cfcPC_\star := (cfcPC.ce ? cfcPC'_\star : cfcPC_\star)$$

already generated in the CFC PC generation. Additionally, we also have to supply the sequential PC for every CFI, as we will see in the description of the branch checker unit, section 3.11.4. For this we compute

$$cfi4l.inlinePC_\star := \text{inc4}(cfi.PC_\star)$$

with $\text{inc4}(\cdot)$ denoting an increment-by-4 function.

3.8.2 Draining and Switching the Issuing IFQ

The second task of the issue selector environment is to drain the issuing IFQ and initiate a switch, in case that a CFI causing a control flow change has been issued.

The decode / issue environment returns for this purpose the $i_\star.issue$ signals; giving a half-unary encoding of the number of issued instructions. The following signals detects, if a CFI has been issued:

$$cfilssued := \left(\bigvee i_\star.issue \wedge i_\star.cfi \right)$$

If a CFI causing a control flow change has been issued, the issuing IFQ must be switched:

$$iIFQswitch := cfilssued \wedge (cfi.ready ? cfi.cfc : cfi4l.cfc)$$

The issuing IFQ is drained by the number of issuing instructions, supplied in half-unary encoding. We have:

$$ifq_i.read_\star := i_\star.issue \wedge (\text{issueIFQ.idx} = i)$$

Decode / Issue Stage

3.9 Decode / Issue Environment

The purpose of the decode / issue environment (DIenv) is to take instructions from the IFQ in a sequential order, compute control signals for them, gather their operands if possible and direct them to the appropriate reservation stations to initiate their execution. These operations have to be executed for an arbitrary but fixed number β of instructions simultaneously.

In accordance to the nomenclature of issue protocol, algorithm 1 (p. 16), the instructions in the instruction window are delivered by the instruction fetch stage on busses $i_i.\star$. The following two subsections describe the exact operation of the DIenv.

3.9.1 Decoding the Instruction

The decoding component for instruction i_i takes the instruction opcode $i_i.opc$ and computes the following control signals:

- Signal $isCFI$ indicates that the decoded instruction is a CFI. The opposite condition, i.e. that the instruction is a BBI, is expressed by the negation, \overline{isCFI} .
- The signals $itype$, $jtype$ and $rtype$ determine the *instruction type* according to the instruction set architecture of the DLX, appendix A. The instruction type specifies the location of operand addresses or immediate constants in the instruction's opcode.

Naturally, exactly one of the signals $itype$, $jtype$ and $rtype$ should be activated.

- Floating point instructions are marked by the signal fp ; in case that they operate on double IEEE encodings, signal db is activated additionally.
- The functional unit identifiers $FU.\star$, explicitly listed in table 3.6, specify which functional unit will be used for the execution of instruction i_i . The special signal $noFU$ indicates, that the instruction does not need a functional unit for execution, i.e. it completes on issue. Only illegal instructions meet this requirement.

Naturally, exactly one of the signals $FU.\star$ should be activated.

- The *operand signals* are defined for the source operands op_1 and op_2 and for the destination operand d of an instruction i_i . Their signals can be divided into two subgroups.

The first group, consisting of the signals IMM , $RS1$, $RS2$, RD , $FS1$, $FS2$, SA , $R31$, describe the possible sources and destination in a DLX instruction. These are, in order of appearance, a 16-bit or 26-bit immediate constant, the first and second general purpose register operand, the integer destination register, the first and second floating point source operand, the shift amount and register R_{31} . Again, the nomenclature and semantics is shown in appendix A.

The second group locates register operands in the machine: gpr , fpr and spr determine if the register is general-purpose, floating-point or special; db specifies a double operand in case of a floating-point reference. The register address is specified by the address bus $a_\star \in \{0, 1\}^5$.

Functional Unit	Purpose
$FU_0 \equiv FU.alu$	int computation
$FU_1 \equiv FU.mem$	load / store
$FU_2 \equiv FU.add$	fp addition / subtraction
$FU_3 \equiv FU.fmul$	fp multiplication
$FU_4 \equiv FU.fdiv$	fp division
$FU_5 \equiv FU.fconv$	conversion int / fp
$FU_6 \equiv FU.ftest$	fp condition test
$FU_7 \equiv FU.bcu$	branch checker unit

Table 3.6: Coding of the functional units

- The remaining signals, load, fabsneg, ff2i, fi2f, fmov, link, jump, noChk, trap, rfe, movs2i, movi2s and ill, specify behaviour of the functional unit executing the instructions. They are just passed on to the functional unit and will not be described in detail here.

The decoding of the instruction, i.e. the computation of the signals just described, is divided into two parts. The control automaton CSig decodes the instruction word and activates the control signals necessary for the execution of the instruction. This is done by specifying states for sets of instructions, as listed in table 3.7. Table 3.8 defines the automaton CSig. For each state, monomials are listed matching to the corresponding instructions (cf. appendix A). The automaton computes the monomials to determine its state at the beginning of each cycle. Then, it activates the control signals for each state, in correspondance to the signals given in the table.

A discussion of such automata, including analysis for cost and delay, is found in [MP95].

Additionally, for each operand op_1 , op_2 and dest the operand address is generated. The operand address is gathered from opcode bits specified in appendix A according to the operand type. So, we have:

$$\begin{aligned}
op_1.a_\star &:= \begin{cases} opc_{10..6} & \text{for } op_1.SA = 1 \\ opc_{20..16} & \text{for } op_1.RS2 = 1 \\ opc_{25..21} & \text{otherwise} \end{cases} \\
op_2.a_\star &:= opc_{20..16} \\
d.a_\star &:= \begin{cases} (11111) & \text{for } d.R31 = 1 \\ (01000) & \text{for } d.FCC = 1 \\ opc_{10..6} & \text{for } d.SA = 1 \\ opc_{15..11} & \text{for } \overline{itype} \wedge (d.RD \vee d.FD) = 1 \\ opc_{20..16} & \text{for } itype \wedge (d.RD \vee d.FD) = 1 \end{cases}
\end{aligned}$$

The computation of the signals db, fpr, gpr and spr for the operands is straightforward. Refer to figure 3.9 showing the operand address computation.

3.9.2 Issuing the Instruction

Testing for Reservation Station Availability

The issue protocol, algorithm 1 (p. 16), states that an instruction i_i cannot be issued, if there is no appropriate reservation station available. This section derives the computation of the signal $i_i.noRS$ indicating this condition.

Consider a functional unit FU_j , fixed. This functional unit is associated with a set of reservation stations which we will later organize in a queue RSQ_j (section 3.14)

State	Instructions
ALU	sll, sla, srl, sra add, addu, sub, subu, and, or, xor, lhg clr, sgr, seq, sge, sls, sne, sle, set
Shifti	slli, slai, srli, srai
ALUi	addi, addiu, subi, subiu andi, ori, xori, lhgi clri, sgri, seqi, sgei, slsi, snei, slei, seti
Load	lb, lh, lw, lbu, lhu
Load.s	load.s
Load.d	load.d
Store	sb, sh, sw
Store.s	store.s
Store.d	store.d
Faddsub.s	fadd.s, fsub.s
Faddsub.d	fadd.d, fsub.d
Fmul.s	fmul.s
Fmul.d	fmul.d
Fdiv.s	fdiv.s
Fdiv.d	fdiv.d
Fcond.s	fc.cond.s
Fcond.d	fc.cond.d
Fabsneg.s	fabs.s, fneg.s
Fabsneg.d	fabs.d, fneg.d
Ff2i	mf2i
Fi2f	mi2f
FMov.s	mov.s
FMov.d	mov.d
FConv.s	cvt.s.d, cvt.s.i, cvt.i.s, cvt.i.d
FConv.d	cvt.d.i, cvt.d.s
Branch	beqz, bnez
FBranch	fbeqz, fbnez
JumpReg	jr
JLinkReg	jalr
Jump	j
JLink	jal
Trap	trap
RFE	rfe
Movs2i	movs2i
Movi2s	movi.s
uFOP	fsqt.s, fsqt.d, frem.s, frem.d

Table 3.7: Correspondance of states and instructions

State	Monomials		Active Control Signals			
	opc _{31..26}	opc ₆ opc _{5..0}		FU.	op ₁ .	op ₂ . d.
ALU	000000	* 0001**	rtype	alu	RS1	RS2 RD
Shifti	000000	* 10****	rtype	alu	RS1	imm RD
ALUi	0*1***	* *****	itype	alu	RS1	imm RD
Load	100***	* *****	itype, load	mem	RS1	– RD
Load.s	110001	* *****	itype, load, fp	mem	RS1	– FD
Load.d	110101	* *****	itype, load, fp, db	mem	RS1	– FD
Store	101***	* *****	itype	mem	RS1	RD –
Store.s	111001	* *****	itype, fp	mem	RS1	FD –
Store.d	111101	* *****	itype, fp, db	mem	RS1	FD –
Faddsub.s	010001	0 00000*	rtype, fp	fadd	FS1	FS2 FD
Faddsub.d	010001	1 00000*	rtype, fp, db	fadd	FS1	FS2 FD
Fmul.s	010001	0 000010	rtype, fp	fmul	FS1	FS2 FD
Fmul.d	010001	1 000010	rtype, fp, db	fmul	FS1	FS2 FD
Fdiv.s	010001	0 000011	rtype, fp	fdiv	FS1	FS2 FD
Fdiv.d	010001	1 000011	rtype, fp, db	fdiv	FS1	FS2 FD
Fcond.s	010001	0 11****	rtype, fp	ftest	FS1	FS2 FCC
Fcond.d	010001	1 11****	rtype, fp, db	ftest	FS1	FS2 FCC
Fabsneg.s	010001	0 00010*	rtype, fabsneg, fp	fconv	FS1	– FD
Fabsneg.d	010001	1 00010*	rtype, fabsneg, fp, db	fconv	FS1	– FD
Ff2i	010001	* 001001	rtype, ff2i, fp	fconv	FS1	– RS2
Fi2f	010001	* 001010	rtype, fi2f, fp	fconv	RS2	– FS1
FMov.s	010001	0 001000	rtype, fmov, fp	fconv	FS1	– FD
FMov.d	010001	1 001000	rtype, fmov, fp, db	fconv	FS1	– FD
FConv.s	010001	* 100*00	rtype, fp	fconv	FS1	– FD
FConv.d	010001	* 100001	rtype, fp, db	fconv	FS1	– FD
Branch	00010*	* *****	itype, isCFl	bcu	RS1	– –
FBranch	00011*	* *****	itype, isCFl	bcu	FCC	– –
JumpReg	010110	* *****	itype, jump, isCFl	bcu	RS1	– –
JLinkReg	010111	* *****	itype, jump, isCFl	bcu	RS1	– R31
Jump	000010	* *****	jtype, jump, noChk, isCFl	bcu	imm	– –
JLink	000011	* *****	jtype, jump, noChk, isCFl	bcu	imm	– R31
Trap	111110	* 000000	jtype, trap, noChk	bcu	imm	– –
RFE	111111	* *****	jtype, rfe, noChk, isCFl	bcu	imm	– –
Movs2i	000000	* 010000	rtype, movs2i	alu	SA	– RD
Movi2s	000000	* 010001	rtype, movi2s	alu	RS1	– SA
uFOP	010001	* 00011*	uFOP	noFU	–	– –
	010001	* 01****				
Illegal (z ₀)	(alternative case)		ill	noFU	–	– –

Table 3.8: Decode control

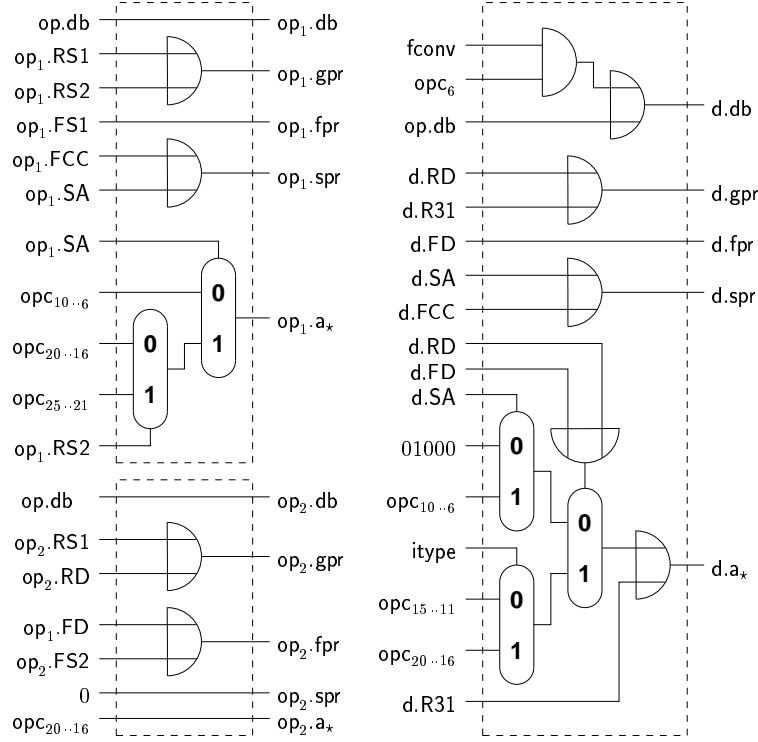


Figure 3.9: Operand address generation

of size $RSQ.SZ_j$. The fill state of the queue is indicated by the signals $RSQ_j.empty_*$; we have $\langle RSQ_j.empty_* \rangle_{hu} = l$ iff there no more than $(l + 1)$ free reservation stations left. Furthermore, the number of instructions issuing on the queue each cycle is bounded by its “indegree”, the parameter $k := RSQ.IN_j$.

The instruction i_i cannot issue on RSQ_j if it is the l -th candidate requesting FU_j and RSQ_j has no l free reservation stations left. In circuitry, we compute the candidate number $i2RSQ_{i,j}.cand_* \in \{0, 1\}^{k+1}$ for instruction i_i issuing on RSQ_j :

$$\langle i2RSQ_{i,j}.cand_* \rangle_{hu} = \min \{k + 1, \text{ones}(i_{i..0}.FU_j)\}$$

A find-first- $(k + 1)$ -ones-half-unary circuit, described in section C.2, computes this information. Now we check $\langle i2RSQ_{i,j}.cand_* \rangle_{hu} > \langle RSQ_j.empty_* \rangle_{hu}$ by half-unary comparison (equation 3.12, p. 36) and obtain

$$RSQbusy4l_{j,i} := \left(\bigvee \overline{i2RSQ_{i,j}.cand_*} \wedge RSQ_j.empty_* \right)$$

The following equation computes $i_i.noRS$:

$$i_i.noRS := \overline{i_i.noFU} \wedge \left(\bigvee i_i.FU_* \wedge RSQbusy4l_{*,i} \right)$$

This reflects the fact that an instruction needing no functional unit never fails in finding one and for $i_i.FU_j = 1$ there must be a free reservation station in RSQ_j , i.e. $RSQbusy4l_{j,i} = 1$.

Source Operand Data Generation

The source operand data generation computes for each source operand a tuple $(tag_*, valid, data_*) \in \{0, 1\}^\nu \times \{0, 1\} \times \{0, 1\}^{32}$, where ν is the tagwidth (cf. section

3.13). Double source operands are composed of a low part and a high part, so four operands $\text{op}_{k,b}$, $k \in \{1, 2\}$ and $b \in \{lo, hi\}$ have to be generated.

The source operand data generation is a direct implementation from the appropriate portion of the issue protocol, algorithm 1, p. 16. Recall that $\text{valid} = 1$ signals a valid data_\star component while $\text{valid} = 0$ ensures that tag_\star contains the tag of the producing instruction of the operand.

We proceed with a description of the various cases. They are ordered according to their priority, so the latter cases only apply if none of the previous cases applies.

- Let $b = lo$. If the operand is the immediate constant, signalled by $i.\text{op}_k.\text{imm}$, we set

$$(\text{tag}_\star, \text{valid}, \text{data}_\star) := (0^\nu, 1, i.\text{co}_\star)$$

To compute the immediate constant $i.\text{co}_\star \in \{0, 1\}^{32}$, three cases must be distinguished: for register type instructions, the immediate constant is taken from the shift amount field, $i.\text{opc}_{10..6}$. Otherwise, it is the sign-extended 16-bit constant for immediate type instructions or the sign-extended 26-bit constant for jump type instructions. The distinction of the three cases results in the following equations:

$$\begin{aligned} i.\text{co}_{5..0} &:= (i.\text{rtype} ? i.\text{opc}_{10..6} : i.\text{opc}_{5..0}) \\ i.\text{co}_{15..6} &:= i.\text{opc}_{15..6} \\ i.\text{co}_{25..16} &:= (i.\text{jtype} ? i.\text{opc}_{25..16} : i.\text{opc}_{15}) \\ i.\text{co}_{32..26} &:= i.\text{co}_{25} \end{aligned}$$

For $b = hi$, the procedure is simplified. If an immediate constant operand is indicated by $i.\text{op}_k.\text{imm}$ or if the operand is no double operand ($i.\text{op}_k.\text{db} = 1$) we define:

$$(\text{tag}_\star, \text{valid}, \text{data}_\star) := (0^\nu, 1, 0^{32})$$

- We now check, if an accompanying instruction i_j with $j < i$ produces the source operand. This is done by comparing the source operand's location information db , gpr , fpr , spr and a_\star to the destination location for i_j . Let the predicate $\text{eqA}(i, k, b, j)$ indicate that operand $i.\text{op}_{k,b}$ is produced by i_j . For $b = lo$ this is the case, if the operand has the same address for the same location. If the destination of i_j is a double floating point register, only the upper four bits of the address must match. So, with S denoting $i.\text{op}_{k,b}$ and D denoting $i_j.d$, we define:

$$\begin{aligned} \text{eqA}(i, k, lo, j) &:= ((S.\text{gpr}, S.\text{fpr}, S.\text{spr}, S.\text{a}_{4..1}, S.\text{a}_0 \vee D.\text{db}) \\ &= (D.\text{gpr}, D.\text{fpr}, D.\text{spr}, D.\text{a}_{4..1}, D.\text{a}_0 \vee D.\text{db})) \end{aligned}$$

For $b = hi$, the computation is similar. This case is only relevant for $i.\text{op}_k$ being a double floating point register, since the other operands are all 32 bits. Instruction i_j produces the source operand, if its destination is a single floating point register with $\text{a}_0 = 1$ or double floating point register. So, with S denoting $i.\text{op}_{k,b}$ and D denoting $i_j.d$, we define:

$$\text{eqA}(i, k, hi, j) := (i.\text{fpr}, \text{a}_{4..1}, 1) = (i_j.\text{fpr}, \text{a}_{4..1}, \text{a}_0 \vee i_j.\text{db})$$

In case that equal (or as in the double case rather overlapping) location is signalled, we obtain the instruction tag $i_j.\text{tag}_\star := \text{ROB.tail}_{j,\star}$ for the maximum j with $\text{eqA}(i, k, b, j) = 1$

$$(\text{tag}_\star, \text{valid}, \text{data}_\star) := (i_j.\text{tag}_\star, 1, 0^{32})$$

- Now we have to access the register file to obtain the tag and the valid bit the source operand. In our implementation, tag and valid bit are stored in the so-called producer tables, while the actual data of registers is stored in the register file.

The producer table is accessed via the addresses of the source operands and returns the signal $PTvalid := pt.i_i.op_{k,b}.valid$ and bus $PTtag_* := pt.i_i.op_{k,b}.tag_*$. If $PTvalid = 1$, then the data stored in the register file, $rf.i_i.op_b.data_{b,*}$ is valid and can be used to compose the source operand:

$$(tag_*, valid, data_*) := (0^\nu, 1, rf.i_i.op_i.data_{b,*})$$

- Having $PTvalid = 0$, the instruction producing the source operand has not yet retired. To detect the result being broadcast on a common data bus in the same cycle, we define

$$snoopCDB_1 := ((CDB_1.valid, CDB_1.tag_*) = (1, PTtag_*))$$

We have $snoopCDB_1 = 1$, if the result has been detected on CDB_1 . Note, that as tags are unique, at most one CDB can qualify in having the tag being looked for, so the order of checks is not important. A successful snooping has taken place, if one of the signals $snoopCDB_1$ is active:

$$snoop := \left(\bigvee snoopCDB_1 \right)$$

In case of snooping, data from the snooped CDB can be driven by $snoopCDB_1$ on the source operand bus.

- If the reorder buffer, accessed by $PTtag_*$, signals valid data with $rob.i_i.op_i.valid$, the instruction having tag $PTtag_*$ already completed its computation and the result is stored in the reorder buffer as $rob.i_i.op_{k,b}.data_*$. So, under the assumption of $rob.i_i.op_{k,b}.valid = 1$, we set:

$$(tag_*, valid, data_*) := (0^\nu, 1, rob.i_i.op_{k,b}.data_*)$$

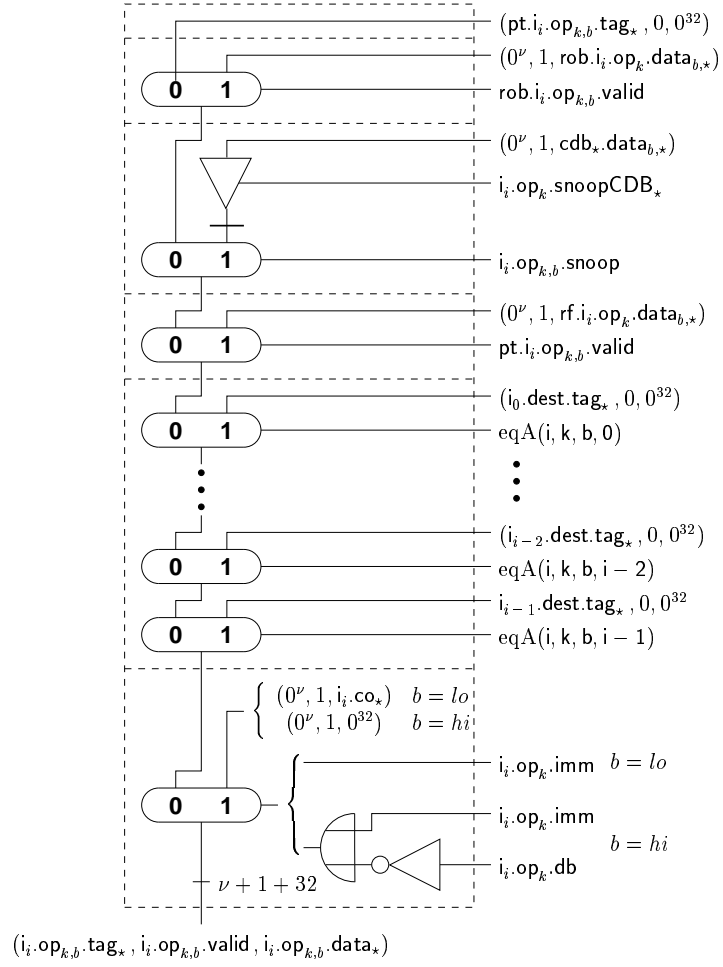
- Finally, only the tag is available in this case. Therefore we set

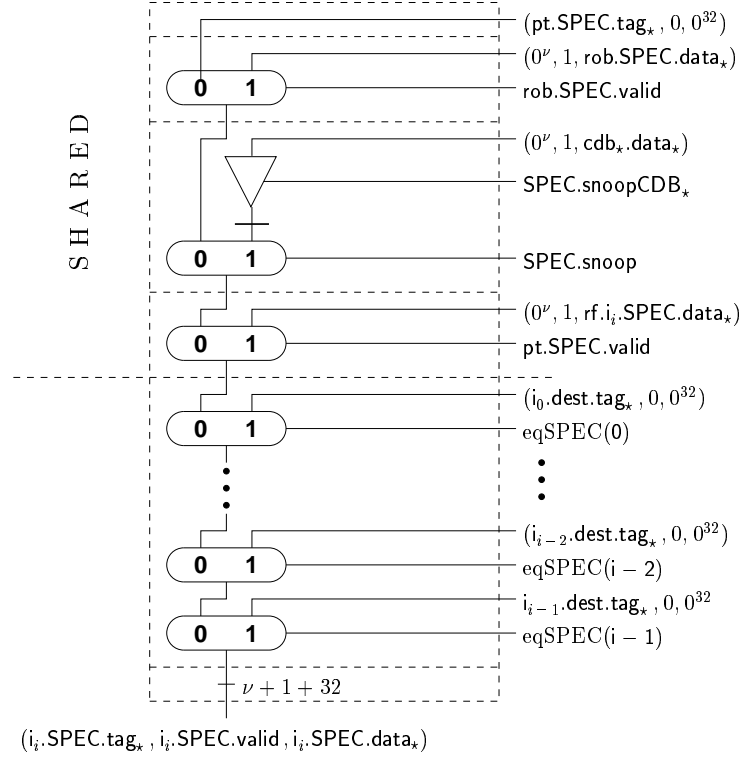
$$(tag_*, valid, data_*) := (PTtag_*, 0, 0^{32})$$

The cases just described lead to the implementation of the source operand data generation as shown in figure 3.10.

Figure 3.11 shows the source operand data generation for the special registers RM (rounding mode) and MSK (interrupt mask) needed by floating point operations. The procedure is similar to the general operands. Since these source operand generation refers to fixed-address registers, most part of its hardware may be shared for different instructions. The predicates $eqRM(j)$ and $eqMSK(j)$ are used to detect instructions writing to RM and MSK. With the addresses of the RM and the MSK register being $\langle 0110 \rangle_2$ and $\langle 0000 \rangle_2$, according to table 3.13, these predicates are defined as follows:

$$\begin{aligned} eqRM(j) &:= (i_j.d.spr, i_j.a_{3..0}) = (1, 0110) \\ eqMSK(j) &:= (i_j.d.spr, i_j.a_{3..0}) = (1, 0000) \end{aligned}$$

Figure 3.10: Generating the source operand data, $b \in \{lo, hi\}$

Figure 3.11: Generating the source operand data, $SPEC \in \{RM, MSK\}$

Stall Generation

The purpose of the stall generation is to model the execution of the issue loop present in the issue protocol, algorithm 1 (p. 16). The goal is the computation of the signal $i_i.stall$ (and its complement $i_i.issue := \overline{i_i.stall}$); $i_i.stall = 1$ iff instruction i_i cannot be issued. We proceed in two steps. First, we define the signal $i_i.stallAux$, indicating a “local” stall condition. This signal is used to compute the actual stall signals in the second step.

The following cases form the local stall condition for an instruction i_i :

- On rollback condition, signalled by `rollback`, no instruction shall issue.
- The cycle after an `rfe` instruction has been issued, the machine performs the actual register transfers for returing from exception (i.e. `SR` will be initialized to 0). The register `doRFE` is used to catch this event and stall issuing in this case:

$$doRFE' := \left(\bigvee i_*.rfe \wedge i_*.issue \right)$$

- The instruction may not be ready for issue, which happens in two cases: the instruction is not fetched or the instruction was falsely classified by the instruction fetch mechanism. This results in

$$i_i.invalid := \overline{i_i.valid} \vee (i_i.cfi \oplus i_i.isCFI) \vee (i_i.bbi \oplus \overline{i_i.isCFI})$$

- There is no reservation station left to issue instruction i_i , i.e. $i_i.noRS = 1$.
- There is no room left to store i_i in the reorder buffer. The reorder buffer signals this condition by $ROB.i_i.full = 1$.




Figure 3.12: Eligibility for read-out for reservation station

- The instruction is a `movs2i` with source register `IEEEf*` and the ROB is not empty. This is necessary since no forwarding mechanism for the `IEEEf*` register is implemented, so preceding floating point instructions may modify `IEEEf*` without notice. This condition is written as follows:

$$i_i.\text{IEEEfstall} := i_i.\text{movs2i} \wedge (i_i.\text{opc}_{10..6} = 00111) \wedge \overline{\text{ROB.empty}}$$

The following equation summarizes all the above cases:

$$\begin{aligned} i_i.\text{stallAux} &:= \text{rollback} \vee \text{doRFE} \\ &\vee i_i.\text{invalid} \vee i_i.\text{noRS} \vee \text{ROB}.i_i.\text{full} \vee i_i.\text{IEEEfstall} \end{aligned}$$

The (global) stall condition is written as follows:

$$i_i.\text{stall} := \left(\bigvee i_{i..0}.\text{stallAux} \right)$$

This information can be computed using a parallel-prefix OR or a find-first-1 half-unary circuit.

3.10 Reservation Station Environment

The Reservation Station Environment is built up using a RS-Queue. An RS-Queue is a queue that provides an additional signal `e4ro` for each entry which is active if the entry is eligible for read-out. Among the candidates for read-out a RS-queue selects the oldest. The formal specification of an RS-queue and its implementation can be found in section 5.4.

The data elements of the RS queue consist of the control signals for the FU, tuples of `(tag, valid, data)` for each source operand and the tag for the destination operand. An entry is eligible for read-out if it is not empty (that is accounted for automatically in the RS queue) and if all its source operands are valid. The resulting circuit is drawn in figure 3.12.

The reservation station queues are initialized on the activation of the global signal `rollback`.

3.10.1 Scheduling of the reservation station input busses

Let RSQ_j have indegree RSQ.IN_j . The instruction destined for FU_j must be mapped on to the reservation stations write busses $\text{RSQ}_j.\text{W}_{*,*}$. We control this mapping by the signal $i2\text{RSQ}2\text{W}_{i,j,k}$; $i2\text{RSQ}2\text{W}_{i,j,k}$ indicates that instruction i_i forwards its operand to the k -th write bus of the RSQ_j .

Instruction i_i is mapped on the k -th write bus of RSQ_j , iff it requests for FU_j and is the $(k+1)$ -th candidate in doing so. Since $i2\text{RSQ}_{i,j}.\text{cand}_*$, defined in section 3.9.2 on page 52, is a half-unary encoding of the candidate number, we have

$$i2\text{RSQ}2\text{W}_{i,j,*} := i_i.\text{FU}_j \wedge \text{markLastOne}(i2\text{RSQ}_{i,j}.\text{cand}_*)$$




Figure 3.13: Common data bus snooping

The interface of the reservation station queues requires us to indicate the write requests. The k -th write bus of RSQ_j is requested, if there exists an instruction actually *issuing* on the k -th write bus of RSQ_j . So:

$$RSQ_j.W.req_k := \left(\bigvee i2RSQ2W_{*,j,k} \wedge i_{*}.issue \right)$$

3.10.2 Snooping for source operands

This section implements the common data bus snooping protocol, algorithm 2, p. 17.

Have a reservation station RS and one of its operands $RS.op_i$. To scan for the presence of the operand's tag on CDB_j we define the signal $RS.op_i.snoopCDB_j$ as follows:

$$\begin{aligned} RS.op_i.snoopCDB_j &:= ((RS_i.op_i.valid, 1, RS_i.op_i.tag_*) \\ &= (0, CDB_j.valid, CDB_j.tag_*)) \end{aligned}$$

The reservation snoops an operand, if one of the signals $RS.op_i.snoopCDB_j$ is active:

$$RS.op_i.snoop := \left(\bigvee RS.op_i.snoopCDB_{*} \right)$$

The following equation is used to update the contents of the $RS.op_i.valid$ register:

$$RS.op_i.valid' := RS.op_i.valid \vee RS.op_i.snoop$$

The definition of these signals justifies the implementation the update circuits for the data fields given in figure 3.13. The signal $snoopCDB_j$ is used as output enable signal to drive $CDB_j.*$ on an internal update bus. The mux selects between the internal update bus and the old data $RS.op_i.*$ with the signal $RS.op_i.snoop$.

Execution Stage

3.11 Function Unit Environments

3.11.1 Arithmetical and Logical Functional Unit

The arithmetical and logical functional unit (ALU) executes the integer compute instructions. For the distinction of these operations, the following items need to be stored in its reservation stations:

$opc_{4..0}$	Operation
00000	left shift
00010	right shift
00011	arithmetic right shift
10000	add without overflow
10001	add with overflow
10010	subtract without overflow
10011	subtract with overflow
10100	bitwise AND
10101	bitwise OR
10110	bitwise XOR
10111	load high
11001	test “>”
11010	test “=”
11011	test “≥”
11100	test “<”
11101	test “≠”
11110	test “≤”

Table 3.9: Encoding of the ALU operations

- The operation code $op_* \in \{0, 1\}^5$. The operation code is gathered from the instruction’s opcode. Let instruction i_i issue on the ALU. Then:

$$\begin{aligned}
op_4 &:= (i_i.type ? 1 : opc_5) \\
op_3 &:= (i_i.type ? opc_{30} : opc_3) \\
op_2 &:= (i_i.type ? opc_{28} : opc_2 \wedge opc_5) \\
op_{1..0} &:= (i_i.type ? opc_{27..26} : opc_{1..0})
\end{aligned}$$

Table 3.9 shows an the encoding of the ALU operations.

- Signals `movi2s` and `movi2s` indicate special move operations.

The design of the ALU is copied from [MP95, MP00] and is therefore not presented here.

3.11.2 Floating Point Functional Units

The machines has five floating point units for IEEE-compliant addition, multiplication, subtraction, relation checking and conversion. We use the same configuration as presented in [Krö99].

The design of the floating point functional units lies beyond the scope of this thesis. The design of the units is treated in [Lei99, MP00] from which they are copied. [Krö99] provides more detail on the integration of the FPUs in Tomasulo design.

3.11.3 Data Memory Functional Unit

The data memory functional unit can be copied from [Krö99]. We only describe here the structure of its reservation station queue, containing the modified superscalar dispatch protocol.

Reservation Station Entry

A reservation station RS_i for the data memory functional unit contains the following items:

- The first operand, contained in $op_1.data_*$, is used as address operand for a memory access. On issue, $op_1.data_*$ is initialized with the sum of the immediate constant and the first operand:

$$op_1.data_* := add(i_1.op_1.data_*, i_1.co_*)$$

This sum equals the immediate constant, if the first operand was not valid on issue. Eventually, this operand will be snooped and then added to $op_1.data_*$. This process is defined below.

- Items $op_2.valid$ and $op_2.data_*$ are the standard fields for the second source operand.
- Item $load$ is active for a load instruction, inactive for a store instruction.
- Item fp signals a floating point instruction.
- Item db signals a double floating point operand.
- The least significant bit of the second operand's address, $op_2.a_0$, is used for the single adjustment of floating point operands: single operands with $op_2.a_0 = 1$ arriving on $op_2.data_{63..32}$ will be shifted to the lower portion of the bus on entering the functional unit.
- The least significant bit of the destination operand's address, $d.a_0$, is likewise used for single adjustment on leaving the functional unit.
- Instruction opcode bits $opc_{28..26}$ encode the operand width and whether the operation is signed or unsigned.

Operand Snooping

We only describe snooping for the address operand. Let RS_i be the i -th reservation station in the data memory reservation station queue. The signal $RS_i.snoopCDB_j$ is defined to detect the operand's tag on CDB_j :

$$\begin{aligned} RS_i.op_1.snoopCDB_j &:= ((RS_i.op_1.valid, 1, RS_i.op_1.tag_*) \\ &= (0, CDB_j.valid, CDB_j.tag_*)) \end{aligned}$$

Operand $RS_i.op_1$ can be snooped, if one of these signals is active:

$$RS_i.op_1.snoop := \left(\bigvee RS_i.op_1.snoopCDB_* \right)$$

The new valid register is computed as:

$$RS_i.op_1.valid := RS_i.op_1.valid \vee RS_i.op_1.snoop$$

The update circuit for the $RS_i.op_1.data_*$ field is given in figure 3.14. The signal $snoopCDB_j$ is used as output enable signal to drive $CDB_j.data_*$ on the input of a 32-bit adder ADD_{32} . The other operand is $RS_i.op_1.data_*$. The adder's result is the memory access address, in case that $snoop$ is valid.

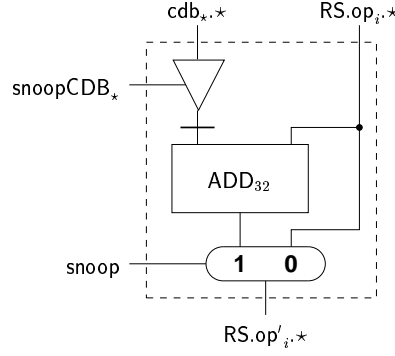


Figure 3.14: Common data bus snooping for the data memory FU

Dispatch Protocol

The interface of the reservation station queue requires the definition of an eligible for read-out signals $e4ro_i$ for each reservation station RS_i . These signals constitute the dispatch protocol.

The dispatch protocol used herein is based on [Mül97]. It is defined separately for load and store instructions:


- A load instruction may dispatch if the following two conditions are met:
 - The operands are valid, i.e. $RS_i.op_1.valid \wedge RS_i.op_2.valid$.
 - Each previous reservation station contains either also a load or its address does not overlap with $RS_i.op_1.data_*$:

$$\forall j < i : RS_j.load \vee (\overline{\text{overlap}(i,j)} \wedge RS_j.op_1.valid)$$

The notion that two addresses overlap is defined by the memory access width. The memory is accessed in 64-bit portions for double operands and in 32-bit portions for all other memory operations. Therefore, if one instruction accesses a 64-bit portion, the accesses overlap, if their upper 29 address bits are the equal. If both instructions are not double, i.e. address 32-bit portions of the memory, their region overlaps, if their upper 30 address bits are equal. These two conditions can be comprised in the following equation of the $\text{overlap}(i,j)$ macro:

$$\begin{aligned} \text{overlap}(i,j) &:= (RS_i.op_1.data_{31..3} = RS_j.op_1.data_{31..3}) \\ &\quad \wedge ((RS_i.op_1.data_2 = RS_j.op_1.data_2) \vee (RS_i.db \vee RS_j.db)) \end{aligned}$$

- Stores may only dispatch, if all preceding instructions have completed. This is the case, if the tag matches the (foremost) head of the ROB queue, $RS_i.tag_* = ROB.head_{0,*}$. This condition implies the following two, noteworthy observations:
 - Due to the organization of reservation station queues, the store instructions is the first in the reservation station queue, since all previous instructions must have completed. Especially, the condition $RS_i.tag_* = ROB.head_{0,*}$ can only be true for $i = 0$ and automatically evaluates to false for $i > 0$.
 - The operands of the store are valid since the preceding instructions have completed and therefore broadcast their results on the common data busses.


$$\begin{aligned} \text{e4ro}_0 &:= (\text{ROB.head}_{0,*} = \text{RS}_0.\text{tag}_*) \\ &\vee (\text{RS}_0.\text{load} \wedge \text{RS}_0.\text{op}_1.\text{valid} \wedge \text{RS}_0.\text{op}_2.\text{valid}) \end{aligned}$$
$$\begin{aligned} \text{e4ro}_i &:= \text{RS}_i.\text{load} \wedge \text{RS}_i.\text{op}_1.\text{valid} \wedge \text{RS}_i.\text{op}_2.\text{valid} \\ &\quad \wedge \bigwedge_{j < i} \text{RS}_j.\text{load} \vee \left(\text{RS}_j.\text{op}_1.\text{valid} \wedge \overline{\text{overlap}(i, j)} \right) \end{aligned}$$

The definition of the rollback protocol (9, p. 28) states, that the predicted next PC has to be compared with the computed next PC to verify a prediction. On a failing verification, the computed next PC has to be reported back to the instruction fetch unit. As has already been outlined in the overview of the implemented instruction fetch mechanism, section 3.3, the verification of CFI predictions is the task of a special functional unit called the branch checker unit, BCU.

A reservation station `RS` of the branch checker unit contains the standard fields `valid`, `op1.valid`, `op1.tag*` and `op1.data*`. In addition, we have the following items to identify control flow instructions and their predictions:

- Signal **jump** indicates the presence of a jump instruction on value 1 and of a branch instruction on value 0.
- Signal **cfc** equals 1, if a CFI was predicted to cause a control flow change. In connection with **jump** = 0, this indicates a taken branch.
- Opcode bit **opc₂₆** is stored to differentiate between the branch types **eqz** and **nez**.
- The signal **noChk** indicates that no verification has to be made for the instruction. Currently, there are three types of instructions just being passed through the BCU without check (cf. 3.8): jump instructions with immediate relative offset, the **trap** instruction and the **rfe** instruction. No speculation is allowed for these instructions, this behaviour may change for branch prediction.
- Signal **trap** indicates the presence of a **trap** instruction in the reservation station. A **trap** instruction passes its immediate constant via **op₁.data_{*}** and the ROB to the **EData** register. This implementation has been chosen to save an additional ROB write port for the issue.
- The **inlinePC_{*}** register contains the inline sequential PC of the instruction.
- The **cfcPC_{*}** register contains the (predicted) PC of a control flow change. In case of a branch instruction, it is the branch target; in case of a jump instruction, it is the predicted PC of a control flow change.

The branch checker unit operates, according to the definition of the DLX ISA, in two different situations:

- On a jump instruction, **jump** = 1, a valid operand **RS.op₁.data_{*}** contains the computed next PC. The predicted next PC is stored in **RS.cfcPC_{*}**. Testing these two values for equality defines the signal

$$\text{bcu.jumpMP} := \left(\bigvee \text{RS.op}_1.\text{data}_*, \text{RS.cfcPC}_* \right)$$

- On a branch instruction, **jump** = 0, we first test the valid operand **op₁.data_{*}** for zero:

$$\text{bcu.op}_1.\text{eqz} := \overline{\left(\bigvee \text{RS.op}_1.\text{data}_* \right)}$$

Since **RS.opc₂₆** = 1 for a **nez** branch and **RS.opc₂₆** = 0 for a **eqz** branch, the correct branch result is computed by

$$\text{bcu.btaken} := \text{bcu.op}_1.\text{eqz} \oplus \text{RS.opc}_{26}$$

The misprediction of a branch is therefore indicated by

$$\text{bcu.branchMP} := \text{RS.cfc} \oplus \text{bcu.btaken}$$

The interface of the branch checker unit is completed with the following definitions:

$$\begin{aligned} \text{bcu.cfcPC}_* &:= (\text{RS.jump} \vee \text{RS.trap} ? \text{RS.op}_1.\text{data}_* \\ &: (\text{bcu.btaken} ? \text{RS.cfcPC}_* : \text{RS.inlinePC}_*)) \end{aligned}$$

$$\text{bcu.mp} := \overline{\text{RS.noChk}} \wedge (\text{RS.jump} ? \text{bcu.jumpMP} : \text{bcu.branchMP})$$

$$\text{bcu.data}_* := \text{RS.inlinePC}_*$$

Refer to figure 3.15 for the actual implementation of the branch checker unit. The computed signals are forwarded to the producer unit and to the branch checker unit bus **bcu.***.

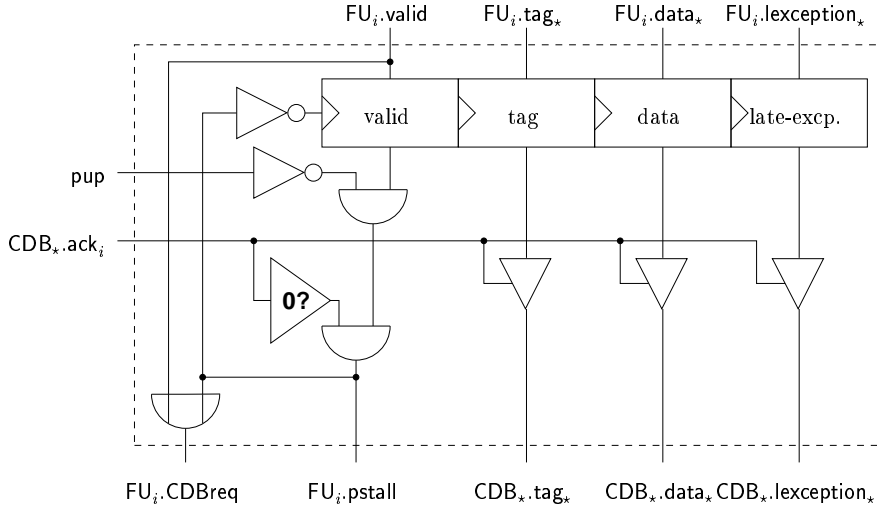


Figure 3.16: Result producer environment

3.11.5 Producer Environments

Implementation

Associated with each functional unit FU_i is a producer environment. The producer buffers results from the functional unit, requests for a common data bus and forwards the result to the CDB on acknowledgement. Figure 3.16 shows the result producer environment for FU_i . The implementation features a central control register *valid*, which operation will now be described.

The functional unit signals a computed result by setting the bit $FU_i.valid$. The producer answers with $FU_i.pstall = 1$ in case that it cannot take the result, because it is already full and did not receive a CDB acknowledgment. With $CDB_j.ack_i$ signalling an acknowledgment from the j -th CDB for producer i , we define:

$$FU_i.pstall := (PROD_i.valid \wedge \overline{pup}) \wedge \left(\bigvee CDB_*.ack_i \right)$$

The *valid* register is clocked, whenever no producer stall is generated:

$$PROD_i.valid.ce := \overline{FU_i.pstall}$$

On clocking the *valid* register, it receives the functional unit's *valid* bit $FU_i.valid$.

A CDB is requested via the signal $FU_i.CDBreq$ in two situation. First, the producer may contain valid data for which it did not receive an acknowledgment. This is the case, if $FU_i.pstall = 1$. Second, a valid FU result requests the CDB. So $FU_i.CDBreq$ is defined as:

$$FU_i.CDBreq := FU_i.valid \vee FU_i.pstall$$

Producer Data Convention

The implementation of our reorder buffer given below requires, that each producer *always* produces 64-bit data. With double IEEE values, this is automatically the case. Single IEEE values, and general purpose or special purpose register result data is 32-bit and needs to be duplicated on the higher 32 bits of the operand bus. So:

$$PROD.data_{63..0} = \begin{cases} FU.data_{63..0} & \text{for 64-bit data} \\ (FU.data_{31..0}, FU.data_{31..0}) & \text{for 32-bit data} \end{cases}$$

3.12 CDB Control Environment

Let k equal the number of common data busses and n equal the number of producers. Already introduced in the producer environment, section 3.11.5, the producer requests the CDB via the signal $FU_{*}.CDBreq$. The task of the CDB control is to fairly select each round up to k producers for acknowledgments. To satisfy the fairness of the selection, we use a k -from- n multiported round-robin selector $k:n$ -MPRRS (cf. section 5.1). Taking $FU_{*}.CDBreq$ as input, this circuit computes the acknowledgment array $ack_{*,*} \in \{0,1\}^{k \cdot n}$ such that the i -th row $ack_{i,*}$ of this array contains a half-unary encoding of the i -th acknowledgment.

So, the signal $CDB_i.ack_j$, acknowledging CDB_i for producer FU_j can be defined via the following equation:

$$CDB_i.ack_{*} := \text{markLastOne}(ack_{i,*})$$

Additionally, the $k:n$ -MPRRS returns the signal $numAck_{*}$, a half-unary encoding of the number of acknowledged candidates. This signal defines the validness of the CDB for the computed acknowledgments:

$$CDB_i.valid := numAck_{*}$$

According to the semantics of the producer bus requests, requesting the CDB for the next round, these signals, $CDB_{*}.ack_{*}$ and $CDB_{*}.valid$, are buffered in flip-flops.

Completion Stage

3.13 Reorder Buffer Environment

The reorder buffer (ROB) guarantees in-order retirement and allows the speculative execution of instruction streams. These properties are necessary to realize precise interrupts and branch prediction. The idea of the reorder buffer was introduced in [SP88].

The ROB presented in this design is based on a multi-ported queue implemented with multiported RAM (section 5.3 describes the implementation of such a queue). In addition to writing entries at the queue tail and reading them at the queue head, the queue must have update ports to read out and write to entries in the “middle” of the queue.

The ROB stores $\Theta = 2^{\nu}$ entries. An entry is a dynamically updated data structure allocated for each issued (but not retired) instruction. The entry stores information about the position, the result and the exception conditions of an instruction. It is referenced by a state-unique id, the instruction tag; a number in the set $\{0, \dots, 2^{\nu} - 1\}$.

Table 3.10 contains a complete list of the ROB entry components. Putting the components into groups provides a way to specify, which item is accessed in which contexts. Of these contexts there are four: During *issue* entries are written in the queue (at the tail). Also during issue, operands are *forwarded* from the ROB to the source operand data generation. Addressing is done via the instruction’s tag, which can be obtained by the producer tables. Third, a *completing* instruction on a CDB must store its results and its exception status in the reorder buffer. Here, the “write address” is given by the instruction’s tag available on the CDB. The last context in which the ROB is accessed is the *retirement* of an instruction. The ROB queue is read out (at its head) for this purpose.

Looking at the components of a ROB entry reveals that they are not accessed in every context. The implementation of the ROB can exploit this information to

Group	Name	Width	Purpose
valid	valid	1	valid data item
dataLow	data _{31..0}	32	low 32 bit of result
dataHigh	data _{63..32}	32	high 32 bit of result
onCompletion	dmal	1	misaligned DMem access
	Dpf	1	data memory page fault
	ovf	1	overflow in ALU instruction
	IEEEf _*	5	IEEE flags
	mp	1	mispredicted CFI
	EData _*	32	exception data
onIssue	ill	1	illegal instructions
	imal	1	misaligned IMem access
	lpf	1	IMem page fault
	trap	1	trap instruction
	uFOP	1	unimplemented fp instruction
	d.a _*	4	destination register address
	d.db	1	double precision
	d.fpr	1	FPR destination
	d.spr	1	SPR destination
	d.gpr	1	GPR destination
	PC _*	32	instruction's PC

Table 3.10: Components of a ROB entry

Group	Width	Ports	Description
valid	1	1I· W: (4I + 2)· RM: 1C· WM: 1R· R:	issue forw. op _{1..2,lo} , op _{1..2,hi} , RM, MSK completion retirement
dataLow	32	(2I + 1)· RM: 1C· WM: 1R· R:	forw. op _{1,lo} , op _{2,lo} , RM completion retirement
dataHigh	32	(2I + 1)· RM: 1C· WM: 1R· R:	forw. op _{1,hi} , op _{2,hi} , MSK completion retirement
onCompletion	41	1C· WM: 1R· R:	completion retirement
onIssue	45	1I· W: 1R· R:	issue retirement

Table 3.11: ROB ports

reduce cost and delay. Critical path analysis in [Krö99] suggests that it may be worthwhile to reduce the delay of the ROB. Table 3.11 contains the of the analysis of the access structure. The table lists for each group the accessing contexts as well as the modes of access (write at tail, read from middle, write in middle, read from head). The constants I , C and R specify the maximum number of issuing, completing and retiring instructions respectively. This information is used for the cost and delay calculation, chapter 4.

3.13.1 ROB Queue Control

The queue is controlled by a number of signals. For queue write operations, the bus $i_*.issue$ supplies the number of instruction to be written in half-unary encoding. The queue returns the address of the i -th entry written on the bus $ROB.tail_{i,*}$. This address is the *tag* of the i -th instruction and is stored in the producer table.

For read operations, the retirement signals $r_*.retire$ generated from the retirement control below are used as read request signals. As shown below, they provide the number of retiring instructions in half-unary encoding; $r_i.retire = 1$ iff the instruction on the i -th retire bus can retire. The tag of the i -th retiring instructions is returned on the bus $ROB.head_{i,*}$.

The read-in-the-middle and the write-in-the-middle ports are directly controlled by appropriate address, request / enable and data-in signals, as specified below.

In addition to this, the queue returns its fill status on bus $ROB.full_*$ that contains a half-unary encoding of the number of occupied entries:

$$i = \# \text{ of occupied entries} \iff ROB.full_* = 0^{\ominus-i} 1^i$$

The queue is initialized on a roll-back condition:

$$ROB.init := rollback$$

3.13.2 Issue

An issuing instruction i_i sets the *valid* and the *onlssue* group of the ROB entry located at the position $ROB.tail_{i,*}$. The components *ill*, *imal*, *IPf*, *trap*, *uFOP*, *d.** and *PC.** are set to their corresponding signals out of the instruction environment. The component *bj* is set to the term $i_i.jumpR \vee i_i.jump \vee i_i.jumpR$. The component *valid* is set to the signal $i_i.noFU$.

3.13.3 Forwarding

In section 3.9.2 the source operand data generation was described. Every instruction i_i requests up to six 32-bit operands to be forwarded from the ROB. These are $op_{1,lo}$, $op_{1,hi}$, $op_{2,lo}$, $op_{2,hi}$, the rounding mode *RM* and the mask register *MSK*. For each operand the data and the valid signals have to be requested.

The producer data convention (section 3.11.5) states, that 32-bit results can be found in both the *dataHigh* and the *dataLow* component. This justifies forwarding $op_{1,lo}$, $op_{2,lo}$ and *RM* from the *dataLow* group and forwarding $op_{1,hi}$, $op_{2,hi}$ and *MSK* from the *dataHigh* group.

Note that without this simple convention, the number of ports is increased for the *dataLow* group resulting in an imbalanced port distribution for *dataLow* and *dataHigh*. There are two reasons for this:

- *Without the convention*, the rounding mode *RM* and the mask register *MSK* would both have to be forwarded from the *dataLow* group. Therefore a read port would have to be removed from *dataHigh* and moved to the *dataLow* group.
- Double floating point source operands may have two producing instructions: one producing the lower 32-bit half of the operand and another one producing the upper 32-bit half of the operand. If the producing instructions are both 32-bit instructions, their result will be stored in the *dataLow* group only *without the convention*. Again, the *dataLow* group would have to provide additional read ports.

3.13.4 Completion

For each valid CDB, the data must be updated in the valid, data and late-exception groups of the corresponding ROB entry. The $CDB_i.valid$ signal provides the write enable signal and the $CDB_i.tag_*$ bus provides the address. The valid component is set to 1. The $dataLow$ and $dataHigh$ components are filled with $CDB_i.data_{31..0}$ and $CDB_i.data_{63..32}$ respectively. The $onCompletion$ group is filled with CDB_i 's late exception signals, $Dmal$, Dpf , ovf , $IEEEf_*$, the misprediction signal mp and the exception data $EData_*$.

3.13.5 Retirement

During retirement exception conditions and branch predictions are checked. Define r_i to be the i -th read-out bus from the ROB. For each r_i the following important interface signals must be generated:

- $r_i.retire$ signals that the instruction retires
- $r_i.wb$ indicates that the result of the instruction actually has to be written back
- $r_i.ilRQ$ signals the detection of an internal interrupt at instruction r_i

Table 3.12 contains a listing of all the possible interrupts. These fall into three categories. The reset exception has the highest priority and is the only unmaskable external interrupt. Exceptions 1 to 12 are the internal exceptions. They are stored in the corresponding ROB entry and are exceptions that occur during the execution of an instruction. The exceptions 13 to 31 are the external I/O exceptions. They are triggered by the external event lines $ex_{18..0}$. The processor detects external interrupt conditions for the *last retiring instruction* (LRI) in a cycle, i.e. the instruction on r_i with $r_i.retire = 1$ and i maximal.

Mispredicted CFIs are handled in much the same way. So, instruction retiring also stops if $r_i.mp = 1$. A mispredicted retiring instruction may only be accompanied by external interrupts or a reset interrupt. All the other interrupts either do not occur in connection with DLX CFIs or their occurrence would have caused the instruction to be transformed in a machine nop (for instruction memory misalignment or instruction memory page fault).

In the following two sections, the computations of control signals for external interrupts and internal interrupts / branch mispredictions is described. Then, in the last section, these results are used to complete the retirement of instructions.

External Interrupts Control Signals

The external interrupt processing requires computing the register $MCA'_{31..13}$ and the signal $elRQ$ signalling an external interrupt request. As external interrupts will be associated with the last retiring instruction lri_* , we define:

$$\begin{aligned} lri.MCA_{31..13} &:= ex_* \wedge SR_{31..13} \\ lri.elRQ &:= \left(\bigvee lri.MCA_{31..13} \right) \end{aligned}$$

Internal Interrupt Control Signals

For each retirement bus r_i a number of signals have to be computed. First, the masked cause bus, $r_i.MCA_*$, is generated, by taking the exception signals and masking them out with the status register SR_* :

Interrupt	Symbol	Priority	Resume	Maskable	External
reset	reset	0	abort	no	yes
illegal instruction	ill	1	abort	no	no
misaligned access	mal	2			
page fault IM	Ipf	3	repeat		
page fault DM	Dpf	4			
trap	trap	5	continue		
FXU overflow	ovf	6	continue	yes	
FPU overflow	fOVF	7	abort / continue		
FPU underflow	fUNF	8			
FPU inexact result	fINX	9			
FPU divide by zero	fDBZ	10			
FPU invalid operation	fINV	11			
FPU unimplemented	uFOP	12	continue		
external I/O	ex _i	12 + i	continue	yes	yes

Table 3.12: Coding of the interrupts

$$\begin{aligned}
r_i.MCA_0 &:= \text{pup} \\
r_i.MCA_1 &:= r_i.\text{ill} \\
r_i.MCA_2 &:= r_i.\text{dmal} \vee r_i.\text{imal} \\
r_i.MCA_3 &:= r_i.\text{lpf} \\
r_i.MCA_4 &:= r_i.\text{Dpf} \\
r_i.MCA_5 &:= r_i.\text{trap} \\
r_i.MCA_6 &:= SR_6 \wedge r_i.\text{ovf} \\
r_i.MCA_{11..7} &:= r_i.IEEEf_* \wedge SR_{11..7} \\
r_i.MCA_{12} &:= r_i.\text{uFOP}
\end{aligned}$$

Second, the following signals are additionally computed for each retiring instruction r_i :

$$\begin{aligned}
r_i.\text{eligible} &:= \text{ROB.full}_i \wedge r_i.\text{valid} \\
r_i.\text{SRdest} &:= r_i.\text{d.spr} \wedge (r_i.\text{d.a}_* = *0000) \\
r_i.\text{ilRQ} &:= \left(\bigvee r_i.MCA_{12..0} \right) \\
r_i.\text{wbAux} &:= \overline{\left(\bigvee r_i.MCA_{4..0} \right)} \vee r_i.\text{mp}
\end{aligned}$$

The meaning of these signals is as follows: $r_i.\text{eligible}$ indicates that there is an instruction at the i -th head of the ROB and it has already completed and may therefore retire; $r_i.\text{SRdest}$ indicates that the instruction writes to the status register SR_* . Only the last retiring instruction in a cycle may write to the status register SR_* , since otherwise interrupts may incorrectly be masked or unmasked. This signal stops instruction retirement beyond r_i . Last, $r_i.\text{ilRQ}$ indicates an internal interrupt at the instruction (an interrupt with priority between 0 and 12). Such an interrupt requires write-back iff $r_i.\text{wbAux}$ is active; mispredicted CFIs will always be written back. If $r_i.\text{ilRQ} = 1$ instruction retirement will also stop.

Completing the Retirement Interface

With the definition

$$r_i.\text{retireVeto} := \overline{r_i.\text{eligible} \vee r_{i-1}.\text{SRdest} \vee r_{i-1}.\text{iIRQ} \vee r_{i-1}.\text{mp}}$$

the $r_i.\text{retire}$ signals can be computed by a parallel-prefix OR:

$$r_i.\text{retire} := \overline{\left(\bigvee_{j=i-1..0} r_j.\text{retireVeto} \right)}$$

The $r_*. \text{retire}$ bus contains the index of the last retiring instruction in half-unary encoding. We convert this information into a unary encoding:

$$r_*. \text{lri} := \text{markLastOne}(r_*. \text{retire})$$

Thereby, $r_i.\text{lri} = 1$ iff the i -th instruction is the last retiring instructions. This signal is used as an output enable signal to construct the bus of the last retiring instruction, $\text{lri}.*$, from the busses $r_*.*$.

Furthermore, we compute the write-back signals $r_*. \text{wb}$; instruction r_i writes back iff it not the last retiring instructions or it is the last retiring instructions and requires write back. This is expressed by:

$$r_i.\text{wb} := r_{i+1}.\text{retire} \vee (r_i.\text{retire} \wedge r_i.\text{wbAux})$$

The global signal **JISR** is set if an external or internal interrupt was found:

$$\text{JISR} := \text{lri}.\text{eIRQ} \vee \text{lri}.\text{iIRQ}$$

The machine performs a rollback on interrupt or misprediction situation, i.e.:

$$\text{rollback} := \text{JISR} \vee \text{lri}.\text{mp}$$

The Write Back Stage

3.14 Register File Environment

The register file environment contains the general purpose register file, the floating point register file and the special purpose register file. The register files are accessed in three different main contexts:

- During issue the source operands are read from the register files if the producer tables indicates their validness. This requires $2I$ read ports since each instruction has two variable-address operands.
- During retirement, the register file environment takes each retirement bus r_j and writes it back to the appropriate register files on activation of $r_j.\text{wb}$. This requires R write ports.
- The special purpose register file has extended modes of access during exception conditions specified in detail in section 3.14.3.

The section proceeds with a description of each register file.

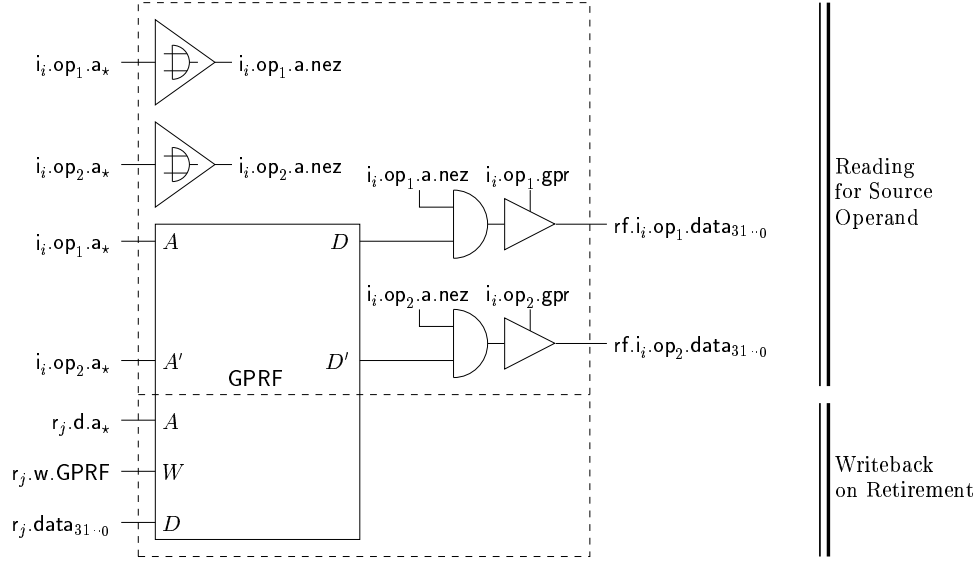


Figure 3.17: Operand reading and retirement ports for the GPRF

3.14.1 General Purpose Register File

The general purpose register file GPRF consists of 32 integer registers having a width of 32 bits. A 32×32 -RAM is used for the implementation.

This RAM has to provide two read ports per instruction for the operands $i_i.op_1$ and $i_i.op_2$. These ports are addressed by $i_i.op_1.a_*$ and $i_i.op_2$ respectively. Since accesses to GPR R_0 must always return 0^{32} , we check if

$$i_i.op_j.a.nez := \left(\bigvee i_i.op_j.a_* \right)$$

and force the output of the RAMs to zero in this case. The data is driven on the operand bus, if GPR data was requested, as indicated by $i_i.op_1.gpr$ and $i_i.op_2.gpr$. The arrangement is shown in figure 3.17.

Additionally, the GPR must also have one write port per retirement bus r_j . This port takes the address $r_j.d.a_*$ and the input data $r_j.data_{31:0}$. The write enable signal $r_j.w.GPRF$ are computed as follows:

$$r_j.w.GPRF := r_j.d.gpr \wedge r_j.wb$$

Again, figure 3.17 shows the arrangement.

3.14.2 Floating Point Register File

The floating point register file consists of 32 single precision floating point registers, alternatively accessed as 16 double precision floating point register. It is constructed of two 16×32 -RAMs, named FPRF0 and FPRF1. The FPRF0 RAM stores the single precision registers with even addresses, the FPRF1 RAM stores the single precision registers with odd addresses. Double precision registers are decomposed in two 32-bit parts; the low part is stored in FPRF0, the high part is stored in FPRF1.

As with the GPRF, two read ports per instruction for the operands $i_i.op_1$ and $i_i.op_2$ are provided. These ports are addressed by $i_i.op_1.a_{4:1}$ and $i_i.op_2.a_{4:1}$ respectively. The data is driven on the operand bus, if a floating point register was requested, as indicated by $i_i.op_1.fpr$ and $i_i.op_2.fpr$. The arrangement is shown in figure 3.18.

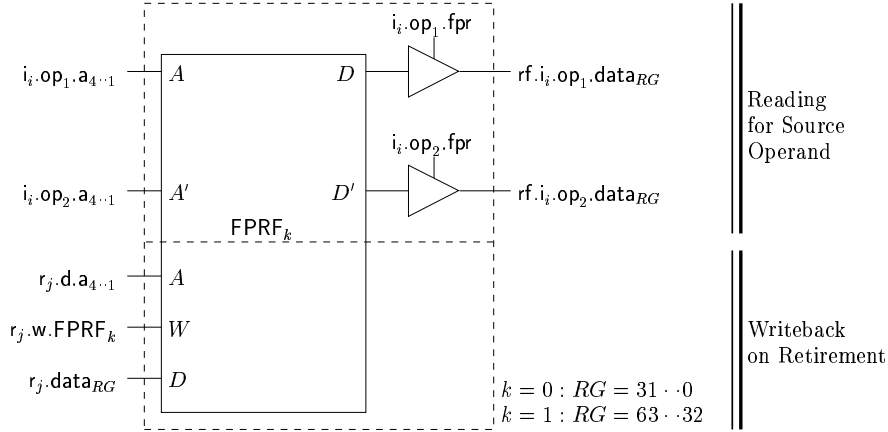


Figure 3.18: Operand reading and retirement ports for the FPRF

Special Purpose Register	Width	Purpose
$\text{SPR}_{0,*} \equiv \text{SR}_*$	32	Status register (interrupt mask)
$\text{SPR}_{1,*} \equiv \text{ESR}_*$	32	Exception status register
$\text{SPR}_{2,*} \equiv \text{EPC}_*$	32	Exception program counter
$\text{SPR}_{3,*} \equiv \text{EPCn}_*$	32	Exception program counter 2 <i>obsolete!</i>
$\text{SPR}_{4,*} \equiv \text{ECA}_*$	32	Exception cause register
$\text{SPR}_{5,*} \equiv \text{EDat}_*$	32	Exception data register
$\text{SPR}_{6,*} \equiv \text{RM}_*$	2	Rounding mode
$\text{SPR}_{7,*} \equiv \text{IEEEf}_*$	5	IEEE interrupt flags
$\text{SPR}_{8,*} \equiv \text{FCC}$	1	Floating point comparison flags

Table 3.13: Special purpose registers

One write port is needed for each retirement bus r_j . The address is provided by $r_j.d.a_{4..1}$. According to the number format indicated by $r_j.db$ and the lowest address bit $r_j.d.a_0$ we define the write enable signals for the FPRi as:

$$r_j.w.FPRF_i := r_j.d.fpr \wedge (r_j.d.db \vee (r_j.d.a_0 = i)) \wedge r_j.wb$$

Again, figure 3.18 shows the arrangement.

3.14.3 Special Purpose Register File

The special purpose register file SPRF consists of nine variable-width registers, as listed in table 3.13. Note that the EPCn_* register is specific for the implementation of a delayed branch. Since our processor does not have a delayed branch mechanism it is not used. Numbering of the special purpose registers has been preserved, though, to ensure address register compatibility to the DLX processors of [MP95, MP00, Krö99, Lei99]. The SPRF is build up using flip-flops for the register storage, since the registers must have extended access structures.

Two read ports per instruction for the operands $i_i.op_1$ and $i_i.op_2$ are provided.⁵ These ports are addressed by $i_i.op_1.a_*$ and $i_i.op_2.a_*$ respectively. The data is driven on the operand bus, if SPRF data was requested, as indicated by $i_i.op_1.spr$ and $i_i.op_2.spr$. If $i_i.op_1.a_{un8..0}$ is a unary encoding of the address $i_i.op_1.a_*$, then the signal

$$i_i.op_1.SPRread_k := i_i.op_1.spr \wedge i_i.op_1.a_{un_k}$$

⁵Instruction access the SPRF only via the first operand, though.




Figure 3.19: Reading from the special purpose register file, $k \in \{0, 1\}$

can be used as an output enable signal to drive special register $\text{SPR}_{k,*}$ on the operand bus $\text{rf.i}_i.\text{op}_1.\text{data}_{31..0}$. The same can be done for $\text{i}_i.\text{op}_2$. The arrangement is shown in figure 3.19.

Additionally, each instruction might request RM and MSK. These requests can be hard-wired though, since they are fixed address and the underlying implementation is based on flip-flops.

For write purposes we always clock each register and consider three modes of access:

- The first mode is the *update* access. This write access is default for all registers. It feeds the data $\text{SPR}_j.\text{upd}_*$ in register SPR_j . Usually we just have

$$\text{SPR}_j.\text{upd}_* = \text{SPR}_j$$

which results in $\text{SPR}'_{j,*} = \text{SPR}_{j,*}$ in case of update access. The exception is the IEEE mask register IEEEf_* , which accumulates the mask of all retiring instructions on normal update:

$$\text{IEEEf}.\text{upd}_* = \bigvee_j (r_j.\text{wb} \wedge r_j.\text{IEEEf}_*)$$

- The second mode of access is the regular write access. Let $r_j.\text{aun}_{8..0}$ be a unary encoding of the retirement address for r_j . Then

$$r_j.\text{SPR}.\text{wreq}_k := r_j.\text{d.spr} \wedge r_j.\text{aun}_k \wedge r_j.\text{wb}$$

indicates a write request for $\text{SPR}_{k,*}$ by retirement bus r_j . The signal

$$\text{SPR}_k.\text{write} := \left(\bigvee_{r_*} r_*. \text{SPR}.\text{wreq}_k \right)$$

detects if $\text{SPR}_{k,*}$ is written to. In this case, the data written must be

$$\text{SPR}_k.\text{writeData}_* = r_j.\text{data} \quad \text{with } j = \max \{j' \mid r_{j'}.\text{SPR}.\text{wreq}_k = 1\}$$

Layers of muxes are sufficient to solve this equation in circuitry.

- The *exceptional* write access refers to the exception status registers $\text{SPR}_{1,*}$ to $\text{SPR}_{5,*}$ only. It is activated on encountering an interrupt situation, $\text{JISR} = 1$.

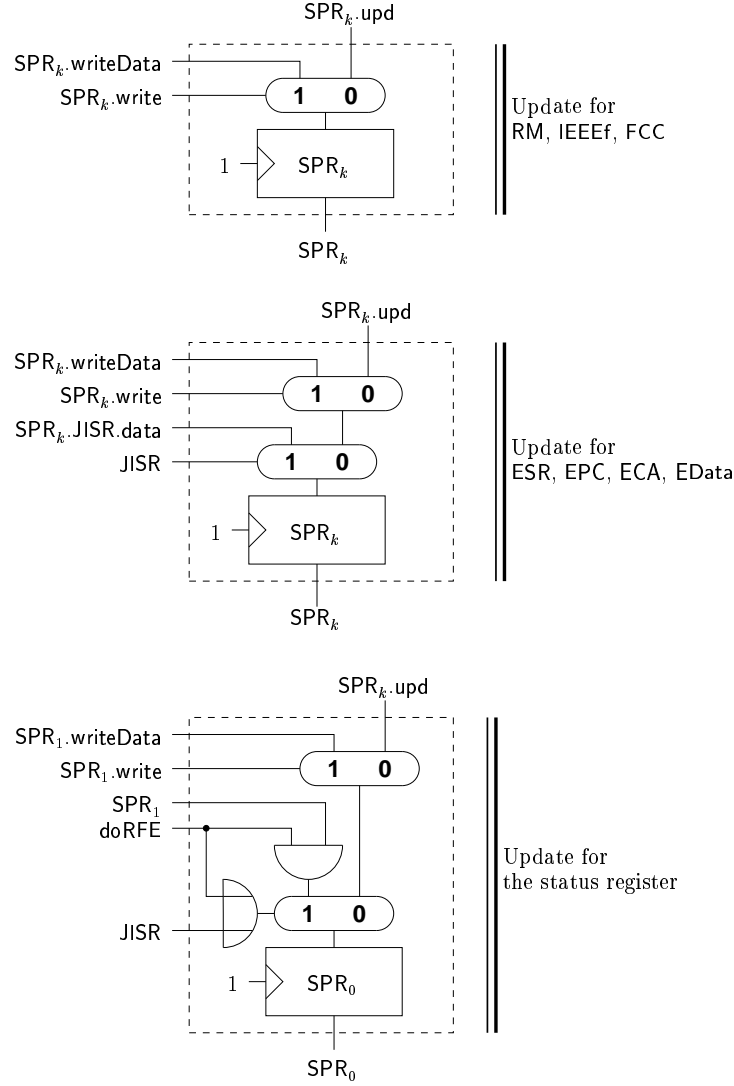


Figure 3.20: Writing to the special purpose register file

The following data is fed into the various special registers:

```

SR.JISR.data* := 0
ESR.JISR.data* := SR*
Iri.nPC* := inc4(Iri.PC*)
EPC.JISR.data* := (Iri.mp ? Iri.EData* : (Iri.wb ? Iri.nPC* : Iri.PC*))
ECA.JISR.data* := Iri.MCA*
EData.JISR.data* := Iri.EData*

```

- Finally, the status registers SR has to be cleared if $doRFE$ is signalled (cf. 3.9, p. 3.9.2 for the definition of $doRFE$).

Figure 3.20 shows the update paths for the special purpose registers.

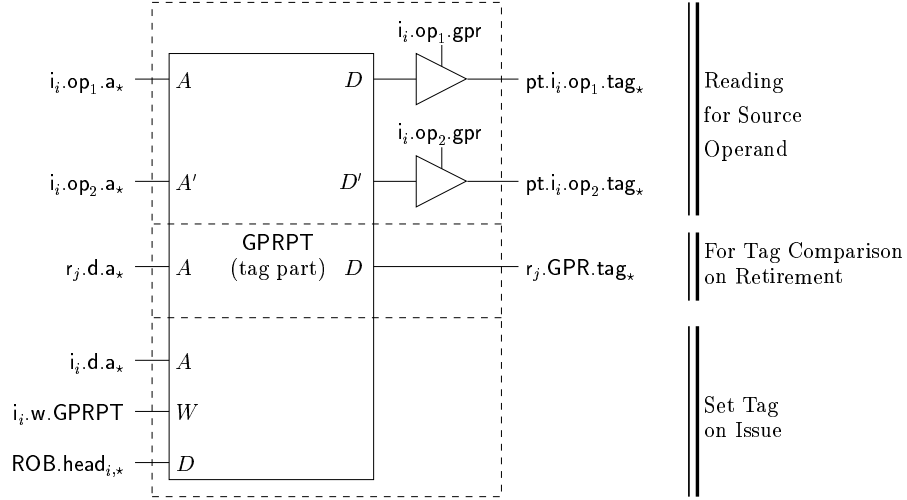


Figure 3.21: Tag information of the GPR producer table

3.15 Producer Table Environment

Producer tables indicate, whether data stored in the register files is valid or being computed in the processor. If the data is still being computed, the producer tables return an identifying tag for the instruction producing the result.

The producer tables are accessed in four contexts:

- First, during issue, information on the valid bit and the tag are taken out of the producer tables for the generation of the source operands $i_i.op_i$.
- Second, the producer tables must be read out during retirement, to determine if a retiring instruction is the producing instruction for its destination register.
- Third, if tags matched in the above context, a retiring instruction stores its result in the register file and may therefore validate the register in the producer table.
- Fourth, destination registers of instructions are marked invalid during issue and a new tag is stored for them. This operation has priority over the third operation (cf. section 2.1.2).

Our processor features three producer tables, one for each register type, which are implemented with register-based RAM to allow write concurrency. The different accessing contexts above suggest—similar to the ROB—a division of each producer table into two parts: one storing the valid information and the other storing the tag information. We follow this approach; however, it will turn out that the producer tables do not lie on a critical path in the processor.

3.15.1 General Purpose Register Producer Table

The general purpose register producer table GPRPT producer table consists of a 32×1 valid RAM and of a $32 \times \nu$ tag RAM; ν is the tag width, cf. section 3.13.

Reading the Source Operands

The tag and the valid part provide two read ports per instruction for the operands $i_i.op_1$ and $i_i.op_2$. These ports are addressed by $i_i.op_1.a_*$ and $i_i.op_2$ respectively. Ac-

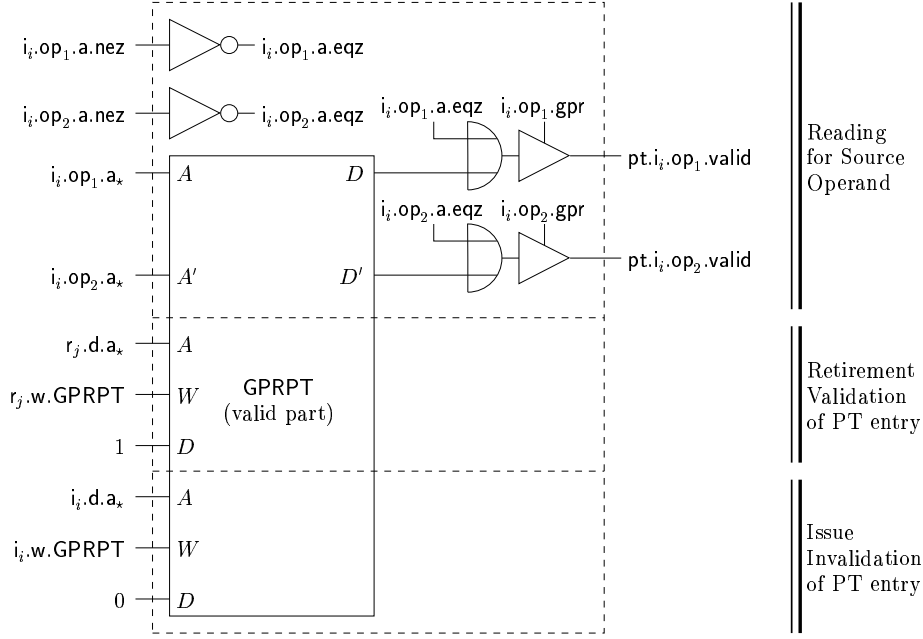


Figure 3.22: Valid information of the GPR producer table

cesses to the GPR R_0 always return a valid flag. We check

$$i_i.op_j.a.eqz := \overline{i_i.op_j.a.nez}$$

and force the output of the valid bit to 1 with an OR gate. The data is only driven on the operand bus, if the instructions are integer operands as indicated by $i_i.op_1.gpr$ and $i_i.op_2.gpr$. Figures 3.21 and 3.22 include this arrangement for the instruction i_i .

Update During Retirement

During retirement, the producer tables tags are read out to decide, if the entry must be validated.

First, the tag of register $r_j.d.a_*$ as stored in the tag part is read out and returned on the bus $r_j.GPR.tag_*$. Define:

$$r_j.w.GPRPT := r_j.wb \wedge r_j.d.gpr \wedge (ROB.head_{j,*} = r_j.GPR.tag_*)$$

This signal equals 1 iff r_j must validate GPR register entry $r_j.d.a_*$ (cf. algorithm 5, page 18, ll. 9–11).

These equations suggest that the full time of a read and a write access to the producer table is needed. However, a closer look the implementation of the register-based RAM reveals that the RAM address decoders for both accesses can work in parallel. This allows for a large reduction in path length.

Figures 3.21 and 3.22 include this arrangement for the retirement bus r_j .

Update during Issue

If instruction i_i issues, it invalidates its destination register in the producer table and sets the tag information accordingly. The signal

$$i_i.w.GPRPT := i_i.issue \wedge i_i.d.gpr$$

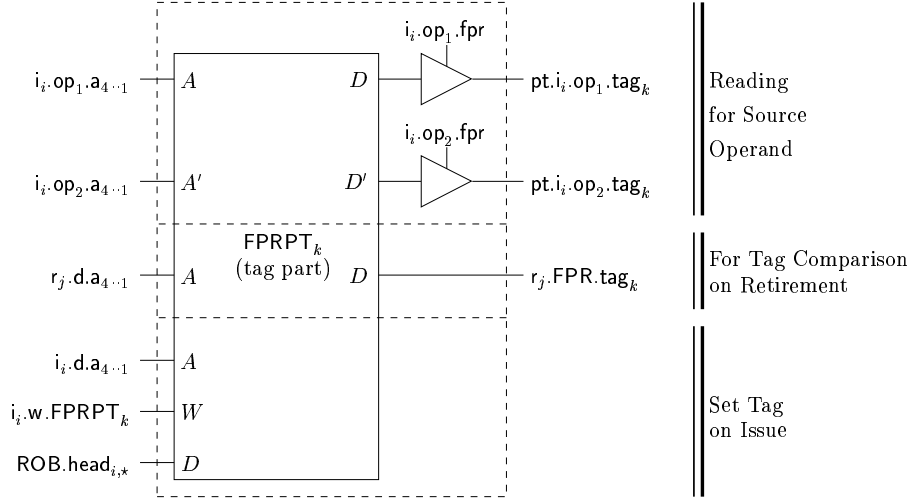


Figure 3.23: Tag information of the FPR producer table

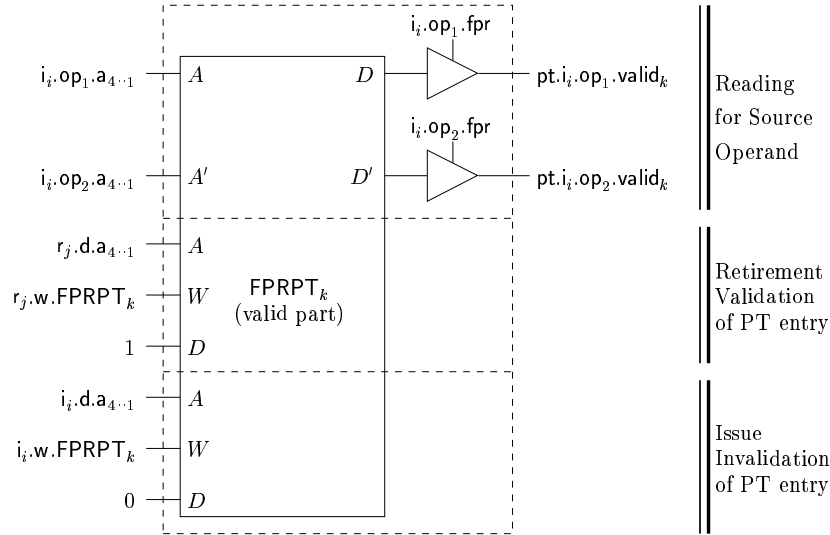


Figure 3.24: Valid information of the FPR producer table

indicates this condition for the GPR producer table. If $i_i.w.GPRPT$ equals one the following two actions are taken:

- The valid part of the GPR producer table stores a 0 at $i_i.d.a_*$.
- The tag part of the GPR producer table stores the instruction's tag, $ROB.tail_{i,*}$, at the address $i_i.d.a_*$.

Figures 3.21 and 3.22 include this arrangement for the instruction i_i .

3.15.2 Floating Point Register Producer Table

The floating point register producer table FPRPT producer table consist of two 16×1 valid RAMs, and of two $16 \times \nu$ tag RAMs. This split is in accordance with the structure of the floating point register file.

Reading the Source Operands

Reading out the tag and valid information for an instruction operand $i_i.op_1$ and $i_i.op_2$ at the operand's addresses $i_i.op_1.a_{4..1}$ and $i_i.op_2.a_{4..1}$ returns two valid flags and two tags for each operand. This data are driven on the operand busses according to the signals $i_i.op_1.fpr$ and $i_i.op_2.fpr$.

Figures 3.23 and 3.24 include this arrangement for the instruction i_i .

Update During Retirement

During retirement, the producer tables tag are read out to determine, if the entry must be validated.

First, we read out the low and high tags of the registers $r_j.d.a_{4..1}$ and return them on the busses $r_j.FPR_{lo}.tag_*$ and $r_j.FPR_{hi}.tag_*$. Then define:

$$\begin{aligned} r_j.w.FPRPT_k &:= r_j.wb \wedge r_j.d.fpr \wedge (r_j.d.db \vee (r_j.d.a_0 = k)) \\ &\quad \wedge (ROB.head_{j,*} = r_j.FPR_k.tag_*) \end{aligned}$$

These two signals indicate, whether the floating point register entry must be validated.

Again, no two real RAM lookups are required for the same reason as stated in 3.15.1, p. 79.

Figures 3.23 and 3.24 include this arrangement for the retirement bus r_j .

Update during Issue

If instruction i_i issues, it invalidates its destination register in the producer table and sets the tag information accordingly. The signals

$$i_i.w.FPRPT_k := i_i.issue \wedge i_i.d.fpr \wedge (i_i.d.db \vee (r_i.d.a_0 = k))$$

indicates this condition for the low and the high part of the FPR producer tables. If $i_i.w.PTGPR_{lo}$ (resp. $i_i.w.PTGPR_{hi}$) equals one the following two actions are taken:

- The valid part of the low FPR producer table (resp. the high FPR producer table) stores a 0 at $i_i.d.a_{4..0}$.
- The tag part of the low FPR producer table (resp. the high FPR producer table) stores the instruction's tag, $ROB.tail_{i,*}$, at the address $i_i.d.a_{4..0}$.

Figures 3.23 and 3.24 include this arrangement for the instruction i_i .

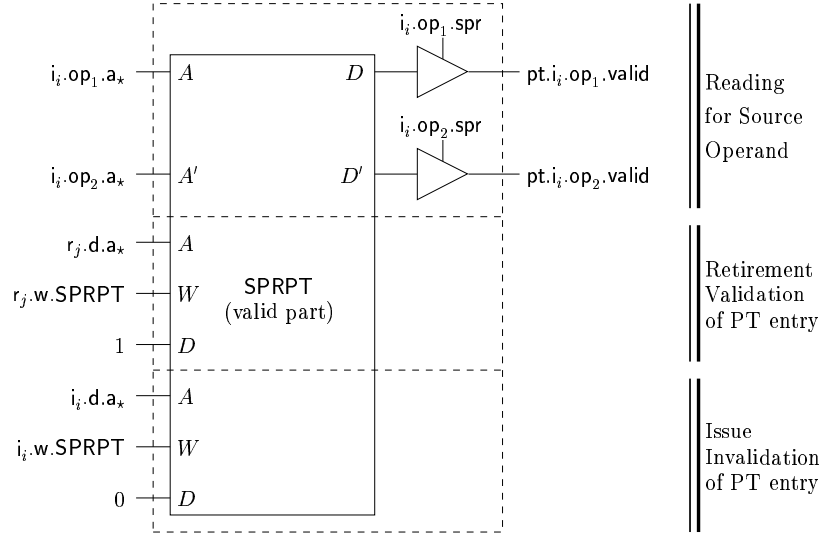
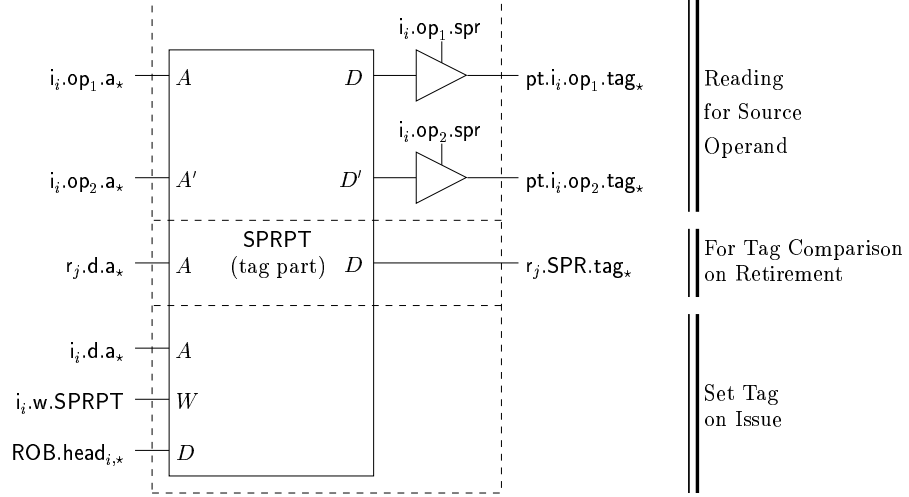
3.15.3 Special Purpose Register Producer Table

The SPR producer table consists of a 32×1 valid RAM and of a $32 \times \nu$ tag RAM.

Reading the Source Operands

The tag and the valid part provide two read ports per instruction for the operands $i_i.op_1$ and $i_i.op_2$. These ports are addressed by $i_i.op_1.a_*$ and $i_i.op_2$ respectively.

Figures 3.25 and 3.26 include this arrangement for the instruction i_i .



Update During Retirement

During retirement, the producer tables tag are read out to determine, if the entry must be validated.

First, the tag of register $r_j.d.a_*$ as stored in the tag part are read out and returned on the bus $r_j.SPR.tag_*$. Define:

$$r_j.w.SPRPT := r_j.wb \wedge r_j.d.spr \wedge (ROB.head_{j,*} = r_j.SPR.tag_*)$$

This signal equals 1 iff r_j must validate the SPRF entry stored at $r_j.d.a_*$.

Figures 3.23 and 3.24 include this arrangement for the retirement bus r_j .

Update during Issue

If instruction i_i issues, it invalidates its destination register in the producer table and sets the tag information accordingly. The signal

$$i_i.w.SPRPT := i_i.issue \wedge i_i.d.spr$$

indicates this condition for the SPR producer table. If $i_i.w.PTSPR$ equals one the following two actions are taken:

- The valid part of the SPR producer table stores a 0 at $i_i.d.a_*$.
- The tag part of the SPR producer table stores the instruction's tag, $ROB.tail_{i,*}$, at the address $i_i.d.a_*$.

Figures 3.25 and 3.26 include this arrangement for the instruction i_i .

Chapter 4

Evaluation

4.1 Hardware Model

The model used herein to determine hardware cost and delay has been introduced in [MP95, KP95]. In this model, basic gates have a fixed cost and delay. The cost of a circuit is defined as the accumulated cost of its gates. The cycle time of a circuit is defined as the longest delay on a path between two registers. Table 4.1 lists the normalized cost and delay for the *MOTOROLA* technology as taken from [MP95].

Note that this model does not take into account fanout restrictions and wire delay. As the Tomasulo algorithm especially in its superscalar variant employs several large bus structures, this topic leaves room for further research.

The cost and the delay of our processor is computed by C++ programs available via WWW¹. They use a circuit library from [MP95] extended by the additional circuits of this thesis. Cost is simply computed by adding up the individual circuits' costs that compose the processor. Delay is computed by modelling the paths of the processor with a C++ class supplied by [Krö99]. This class allows constructing paths by taking the maximum of several paths and by adding delay to a single path. The paths have symbolical names. Two additional programs, both written in perl, were used to simplify the task of constructing the paths. The first program extracts the equations for signal definitions directly from the text of this thesis. This procedure simplifies the error-prone task to maintain the coherence between the delay program and the description of the machine. The second program allows for the modular definition of paths and the instantiation of modules. It takes an arbitrarily ordered input of path equations and sorts these to resolve forward references.

4.2 Parameter Space

The design presented in this thesis leaves a big parameter space to explore. Whenever possible explicit constants have been avoided in the design and have been replaced by a parameter.

The parameters of the processor fall into two classes:

- The *scale parameters* are concerned with the widths of major data paths in the design. They put a theoretical limit on the processor's performance, i.e. the cycles per instruction (CPI) rate.

The parameters we consider belonging to this class, are the number of fetched instructions F , the number of issued instructions I , the number of issues $I(RS)$ on a reservation station RS , the number of dispatches $D(RS)$ of a reservation

¹<http://www-wjp.cs.uni-sb.de/~mah/thesis/costdelay.tgz>

station RS to its functional unit(s), the number of completions C (i.e. the number of CDBs) and the number of retirements R . As a trivial lower bound on the CPI rate we have:

$$CPI \geq \frac{1}{\min\{F, I, C, R\}}$$

- The *size parameters* are concerned with the sizes of the major data structure in our design. Size parameters influence the CPI rate only implicitly: the processor is stalled, if the resources are low. The thought is tempting that by choosing size parameters large enough, stalling can be overcome. However, as we also see below, size parameters have an impact on the processor's cycle time.

The size parameters of our processor that we are concerned with are the size of its instruction fetch queues $IFQSIZE$, the size of the reservation station queues $RSQSIZE(RS)$ for reservation station RS and the size of the reorder buffer $ROBSIZE$. Parameters we leave out of consideration are the BPU table size and the instruction cache size; these parameters do not influence the processor itself.

Because of the lack of simulations, we cannot study the parameter space adequately. We provide a short glimpse on the effects of the scale and the size parameters. For this consider an overall scale parameters sc and a size parameter sz defining the following point in the parameter space:

$$\begin{aligned} F = I = I(ALU) = D(ALU) &= sz \\ I(RS) = D(RS) &= 1 \\ C = R &= sz \\ IFQSIZE &= 8 \cdot sc \\ RSQSIZE(RS) &= 4 \cdot sc \\ RSQSIZE(ALU) &= 4 \cdot sz \cdot sc \\ ROBSIZE &= 2^{4+2(sz-1)} \end{aligned}$$

The choice of parameters reflects the fact that for $sz = sc = 1$, the processor is similar to [Krö99]. Duplication of the ALU with increasing scale sc is suggested by [Joh91] since 40% to 50% of the dynamic instructions of a MIPS instruction set architecture are ALU instructions.

Keeping F , I , C and R to the same value seems to be the natural solution, although it is unclear that this is an optimal choice. This question may be of interest in simulations of the processor.

Tables 4.2 and 4.3 show nine choices for sz and sc . Cost and cycle time grow monotonically with the rows and the columns. These tables must be read with care, since they only concern the machine core. Scaling of the machine has also an impact on components that lie outside the processor core, for example the instruction cache. These must also be adapted, when more power is wished for.

4.3 A 2-Superscalar Processor

We now restrict our choice of parameters further:

- Set $F = 2$, because the underlying instruction cache (cf. [MP00]) has a 64-bit bus to the processor. We assume the modifications necessary to support our instruction fetch protocol are of negligible cost and have no performance impact.

Gate	Cost	Delay
Inverter	$C_{inv} = 1$	$D_{inv} = 1$
NAND	$C_{nand} = 2$	$D_{nand} = 1$
NOR	$C_{nor} = 2$	$D_{nor} = 1$
AND	$C_{and} = 2$	$D_{and} = 2$
OR	$C_{or} = 2$	$D_{or} = 2$
XOR	$C_{xor} = 4$	$D_{xor} = 2$
XNOR	$C_{xnor} = 4$	$D_{xnor} = 2$
Multiplexer	$C_{mux} = 3$	$D_{mux} = 2$
Tristate Driver	$C_{driv} = 5$	$D_{driv} = 2$
Flip-Flop	$C_{ff} = 8$	$D_{ff} = 4$

Table 4.1: Cost and delay of the basic gates

<i>Delay</i>	<i>sz</i> = 1.0		<i>sz</i> = 1.5		<i>sz</i> = 2.0	
<i>sc</i> = 1	107	100.00%	113	105.61%	139	129.91%
<i>sc</i> = 2	122	114.02%	143	133.64%	185	172.90%
<i>sc</i> = 4	164	153.27%	201	187.85%	275	257.01%

Table 4.2: Effect on Delay due to Scaling and Sizing

<i>Cost</i>	<i>sz</i> = 1.0		<i>sz</i> = 1.5		<i>sz</i> = 2.0	
<i>sc</i> = 1	317279	100.00%	424346	133.75%	569332	179.44%
<i>sc</i> = 2	410673	129.44%	551572	173.84%	738716	232.83%
<i>sc</i> = 4	637495	200.93%	863926	272.29%	1153707	363.63%

Table 4.3: Effects on Cost due to Scaling and Sizing

- Consequently, we set $I = C = R = 2$, since there is no hint for a better allocation.
- The ALU is duplicated. Its reservation station are configured for indegree 2 and outdegree 2. So, two instructions may be issued on RS(ALU) and two instructions may also be dispatched from RS(ALU). Studies in [Joh91] suggest that providing an additional ALU in a register-renaming-implementation is worthwhile.
- We have 6 ALU reservations stations and 4 reservation stations for each other functional unit. For a non-superscalar processor 4 and 2 were a good choice according to [Krö99].
- The reorder buffer size is set to 32 to accommodate the additional instructions executed in the superscalar design. Again, 16 seemed to be a wise choice for a non-superscalar processor.
- The IFQ size is set to 8.

4.3.1 Cost and Delay Optimization

Table 4.4 shows a detailed overview of the processor cost.

The execute and the decode / issue stage make up for about 70% of the whole machine cost. In the execute stage, the floating point units are particularly expensive. In the decode / issue stage the reservation station queues contribute most of the cost; note, that there also reservation stations in the DMemFu. So, all the reservation stations cost about half of the machine.

As has been indicated in section 3.13, we consider various variants of the ROB design. First, the ROB may be considered as “one RAM”. This RAM requires $4 * I + 2 + R$ read and $I + C$ write ports. Second, the ROB may be broken down into functional groups, and providing only the ports needed for each group. An overview of the groups is shown in table 3.11. “Trimming” the ROB results in cost savings *and* delay reduction and is therefore the better option in all scenarios.

As we will see, the ROB lies on the critical path of the design. Further delay optimization are therefore advisable. One approach is to replace the regular RAM by register-based RAM. Having the “one RAM” ROB implementation, this results in a fabulous cost increase. With a “trimmed ROB”, however, one is able to optimize the time-critical parts of the ROB only.

Table 4.5 shows a number options. Trimmed ROB is cheapest and has only 50% cost of a regular ROB implementation. Furthermore, the trimmed ROB implementation has only delay 160 compare to gate delay 248 of a regular ROB implementation. Four optimization options are shown for the trimmed ROB. The first one implements the valid group with register-based RAM, the second implements the valid group and the onIssue group with register-based RAM. As these components lie on the critical path, delay decreases from 160 to 143 to 135 with relatively little cost increase. With gate delay 135, the data group lies on the critical path. The data group, divided in low and high data, has a width of 64 bits, and therefore takes up a great portion of the ROB storage. Implementing it with register-based RAM is therefore quite expensive, although delay also reduces considerably.

For comparison to other architectures, we will use the “Trimmed (opt: valid, onIssue)” design with delay 135. Since there the floating point unit has a path of gate delay 137 ([Lei99]), further optimization are not necessary without optimizing the floating point units as well.

Component	Cost		Cf.
	Gates	%	
IF stage (2 fetches, 2 issues)	21349	5.5	
IF Control	385	0.1	3.3.3
IM environment	136	0.0	3.4
IFQ (2 of size 8)	3822	1.0	3.5
PC environment	1563	0.4	3.6
Prediction environment	2682	0.7	3.7
BPU (2-bit sat. cntr, 2048 entries)	8582	2.2	app. B
Instruction Window	357	0.1	3.8
DI stage (2 issues)	111513	28.8	
decode / issue	7678	2.0	3.9
instruction control	2504	0.6	
instruction data	2118	0.5	
global control	410	0.1	
global data	1528	0.4	
reservation stations (snooping from 2 CDBs)	103835	26.9	3.10
EX stage	135658	35.1	
Float Adder	23735	6.1	[Lei99]
Float Mul/Div	47557	12.3	[Lei99]
Float Conv.	15926	4.1	[Lei99]
Float Tranfer	2209	0.6	[Lei99]
Integer ALU (2)	7386	1.9	[MP95]
BCU	461	0.1	3.11.4
DMemFU (including 4 reservation stations)	25664	6.6	[Krö99]
Producers (to 2 CDBs)	12720	3.3	3.11.5
Completion stage (2 compl., 2 ret.)	62940	16.3	
CDB Control	794	0.2	3.12
ROB (32 entries)	61448	15.9	3.13
Retirement Computation / Rollback Checking	698	0.2	3.13.5
WB stage (2 retirements)	55146	14.3	
Register Files	24452	6.3	3.14
Producer Tables	30694	7.9	3.15
Total	386606	100.0	

Table 4.4: Overview of the 2-Superscalar Machine Cost

Type	ROB Cost		Total Cost		Delay
Trimmed	33956	100.00%	359114	100.00%	160
Trimmed (opt: valid)	38132	112.30%	363290	101.16%	143
<i>Trimmed (opt: valid,onIssue)</i>	61448	180.96%	386606	107.66%	135
Trimmed (opt: valid,onIssue,data)	139472	410.74%	464630	129.38%	112
Trimmed (all reg-based)	160716	473.31%	485874	135.30%	112
One RAM	70202	206.74%	395360	110.09%	248
One RAM (reg-based)	428234	1261.14%	753392	209.79%	112

Table 4.5: Comparison of ROB designs

	Pipelined		Tomasulo		Superscalar	
CPU core only	108949	100%	235989	216%	386606	354%
with 16 KB cache	483928	100%	610968	126%	761585	157%
CPI/speedup	2.12	0%	1.47	44%	1.41	50%

Table 4.6: Cost of Complete CPU and CPI rates

0 p_issueIFQ_idx	0 p_issueIFQ_idx
2 issueIFQ[*].*	2 issueIFQ[*].*
0 i_opc	0 i_opc
13 i_CSIG	13 i_CSIG
4 i_op_a	4 i_op_a
9 pt_i_op_tag_GPR	9 pt_i_op_tag_GPR
2 pt_i_op_tag	2 pt_i_op_tag
43 rob_i_op_data	44 rob_i_op_valid
2 i_op_4ROB	2 i_op_4ROB
2 i_op_3CDB	2 i_op_3CDB
2 i_op_2RF	2 i_op_2RF
2 i_op_1DI	2 i_op_1DI
2 cfi_op_validAux	2 cfi_op_validAux
2 cfi_op_valid	2 cfi_op_valid
4 cfi.resolve	4 cfi.resolve
4 cfi.ready	4 cfi.ready
4 i[*].cfi	4 i[*].cfi
6 i[i].invalid	6 i[i].invalid
4 i[i].stallAux	4 i[i].stallAux
2 i[*].stall	2 i[*].stall
1 i[i].issue	1 i[i].issue
4 RSQ[j].W.req[k]	44 rob_issue_VALIDp
0 floatRSQ_r_req	5 p_rob_issue_VALID register in
2 floatRSQ_m	160 TOTAL (23 circuits)
0 floatRSQ_r_ack	
6 floatRSQ_validAfterRead	
1 floatRSQ_w_ack	Trimmed (opt:valid,onIssue)
5 floatRSQ_validAfterWrite	
2 floatRSQ_validP	
0 floatRSQ_data_ce	
5 p_floatRSQ_data register in	
135 TOTAL (31 circuits)	

Trimmed

Table 4.7: Path of maximum delay

4.3.2 Longest Path

Table 4.7 shows a path of maximum delay for two variants of the processor. The processor producing the path on the left-hand side is implemented with a trimmed ROB RAM. The processor producing the path on the right-hand side is the “Trimmed (opt:valid,onIssue)” we have chosen for comparison with other architectures.

The slower, right-hand path starts in the instruction window and continues with instruction decoding. The ROB valid RAM is requested during source operand generation to decide, if the source operand is already stored in the RAM. The constructed source operand is passed on to the fetch mechanism (cfi.★) for resolving branches. Only after the resolving of branches, the instruction fetch mechanism may annotate the instruction window with the bbi and cfi tags and the decode / issue environment may generate the stall signals $i_i.stall$.

The path resembles the critical paths for control flow resolving of classical pipelined designs. It can be avoided with a branch predictor unit predicting *all* types of CFIs, so that CFI resolving will not be necessary.

The faster, left-hand path has the slowest component replaced. However, it takes the same route. The path ends in the reservation station queue control, which can only update its state, if read and write requests have been computed.

4.3.3 Quality Comparison

Quality Measure

This section presents a quality comparison of three DLX designs. Quality is defined, according to [Grü94, MP95], as the q -weighted geometric mean of the performance of an architecture and of the reciprocal of its cost. The weight $q \in [0, 1]$ can be used to put emphasis either on cost or on performance: Only performance counts with $q = 0$, only cost counts with $q = 1$. For $q = 0.5$ cost and performance are equally weighted. The choice of q between 0.2 and 0.5 seems to be realistic.

Justified below, we measure the performance of an architecture by the reciprocal of its CPI rate. Cost is measured by gate counting. Therefore, the quality of a design can be written as follows:

$$Q_q = \frac{1}{CPI^{1-q} C^q}$$

With q fixed, two designs A and B can be compared by their quality Q_q^A and Q_q^B . Better design with respect to q have a greater quality.

Comparison

We compare our design with two others: a pipelined DLX design with precise interrupts and floating point units from [MP95, MP00] and a non-superscalar Tomasulo design presented in [Krö99].

All designs use similar floating point units and therefore share the same longest path with 137 gate delays. Performance therefore depends only on the CPI ratio. The CPI values for the pipelined DLX and the non-superscalar Tomasulo design are taken from [Ger98, Del98]. Due to lack of simulations we have to estimate the CPI rate of the 2-superscalar processor. [Joh91] suggests that a superscalar out-of-order design using 2 ALUs and a *scalar* fetcher has a 1.5-fold speed-up over a non-superscalar processor. We apply this (conservative) estimate to the CPI rate of the pipelined design and obtain a CPI ratio of $2.12/1.5 = 1.41$.

Figure 4.1 plots the quality functions of these designs in a logarithmic scale. The two vertical lines indicate the points of equal quality of the non-superscalar to the

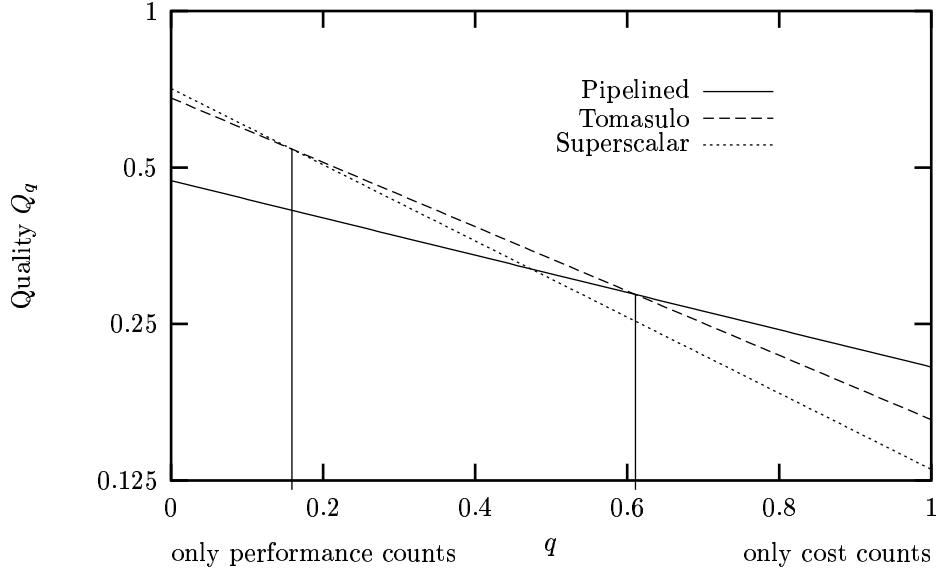


Figure 4.1: Quality for Pipeline DLX, Non-Superscalar and Superscalar Tomasulo

super-scalar and of the non-superscalar to the pipelined processor. With an unrealistic emphasis $q \approx 0.6$ on cost, the pipelined design has the highest quality. For $q \approx 0.18$, the superscalar Tomasulo design wins over the non-superscalar Tomasulo design. This parameter is, however, out of the “realistic” interval. Two reasons can be given for the bad quality: First, our generic approach to superscalar processors most likely produces overhead in cost and delay for instantiation of the machine. This can be mended by optimizing the machine for fixed parameters. Second, simulations miss to replace the CPI ratio 1.41 by a more realistic and hopefully smaller value.

Chapter 5

Circuits

This chapter introduces four circuits forming integral parts of our design.

The multiported round-robin selector presented first, is used to generate bus acknowledgement signals for the common data bus control. Second, a multicounter is developed. This component is used for the implementation of multiported queues in section 5.3. Multiported queues are used in the instruction fetch stage to build the IFQs and in the completion stage to build the reorder buffer. Finally, a (multiported) reservation station queue is developed, providing a multiple-issue and multiple-dispatch framework for reservation stations.

5.1 Multiported Round-Robin Selector

5.1.1 Abstract View on Multiported Round-Robin Selectors

Definition 5.1 Let $n \in \mathbb{N}$ and $N := \{0, \dots, n-1\}$. The round-robin selection function for width n is defined as follows:¹

$$\begin{aligned} \text{rrs}_n : \mathcal{P}(N) \times N &\longrightarrow N \\ (\mathcal{R}, p) &\longmapsto a \end{aligned}$$

such that

$$a = \begin{cases} \min \mathcal{R}^{>p} & \text{if } \mathcal{R}^{>p} \neq \emptyset \\ \min \mathcal{R}^{<p} \cup \{p\} & \text{otherwise} \end{cases}.$$

The function $\text{rrs}(\cdot)$ is used to describe round-robin selecting. \mathcal{R} represents the set of requests received from n busses, p is the previously selected bus and the function result a is the bus that receives acknowledgement.

Next, we define a similar function that handles up to k requests in a round.

Definition 5.2 Let $n \in \mathbb{N}$ and $N := \{0, \dots, n-1\}$. Let $1 \leq k \leq n$ and $K := \{0, \dots, k-1\}$. The k -round-robin selection function for width n is defined as follows:²

$$\begin{aligned} k\text{-rrs}_n : \mathcal{P}(N) \times N &\longrightarrow K \times N^{\leq k} \\ (\mathcal{R}, p) &\longmapsto (\text{numAck}, a_1, \dots, a_{\text{numAck}}) \end{aligned}$$

¹For a set M , the power set is denoted by $\mathcal{P}(M)$. For a relation R on M and $a \in M$ (respectively $b \in M$) define

$$\begin{aligned} M^a R &:= \{b \in M \mid a R b\} \\ M^R b &:= \{a \in M \mid a R b\} \end{aligned}$$

²For a set M define $M^{\leq l}$ as the set of tuples of M with maximum length l . This overloads the notation $M^R b$ defined above; the context, however, always makes the meaning clear.

such that

$$\begin{aligned} numAck &= \min\{k, \#\mathcal{R}\} \\ a_1 &= rrs_n(\mathcal{R}, p) \\ a_i &= rrs_n(\mathcal{R} \setminus \{a_1, \dots, a_{i-1}\}, a_{i-1}) \quad \text{for } i > 1 \end{aligned}$$

The following lemma justifies the implementation of the k -from- n round-robin selector as given below:

Lemma 5.3 *Let $k\text{-}rrs_n(\mathcal{R}, p) = (numAck, a_1, \dots, a_{numAck})$ and let (p_1, \dots, p_l) be the sorted sequence of elements in $\mathcal{R}^{>p}$ and (p_{l+1}, \dots, p_m) be the sorted sequences of elements in $\mathcal{R}^{\leq p}$. Then:*

$$(a_1, \dots, a_{numAck}) \quad \text{is a prefix of} \quad (p_1, \dots, p_m)$$

Proof. We make a finite induction of the number of acknowledgements, $numAck$. For the induction basis, it is clear that a_1 equals p_1 by the definition of $rrs_n(\cdot)$. Assume, that we have showed that (a_1, \dots, a_i) is prefix of (p_1, \dots, p_m) . Then by definition of $k\text{-}rrs_n$ and the induction hypothesis:

$$\begin{aligned} a_{i+1} &= rrs_n(\mathcal{R} \setminus \{a_1, \dots, a_i\}, a_i) \\ &= rrs_n(\{p_{i+1}, \dots, p_m\}, p_i) \\ &= p_{i+1} \end{aligned}$$

5.1.2 Implementation of a Multiported Round-Robin Selector

This section deduces the implementation of a k -from- n round-robin selector. Such a circuit receives n requests signals $\text{req}_\star \in \{0, 1\}^n$; given the encoding of definition 5.2, we have

$$\text{req}_i = 1 \iff i \in \mathcal{R}$$

Furthermore, the circuit stores the index of the last acknowledgement in a register $h_\star \in \{0, 1\}^n$ in (negated) half-unary encoding. Again, with the notation of the definition, we have

$$\langle \overline{h_\star} \rangle_{hu} = p$$

The circuit computes the 2-dimensional signal array $\text{ack}_{\star, \star} \in \{0, 1\}^{k \cdot n}$. The k' -th row $\text{ack}_{k', \star}$ of this array contains a negated half-unary encoding of the k' -th acknowledgement, so

$$\langle \overline{\text{ack}_{k', \star}} \rangle_{hu} = \begin{cases} a_{k'} & \text{if } k' < numAck \\ 0 & \text{otherwise} \end{cases}$$

The register h_\star must be set each cycle to the last acknowledgement i.e.

$$\langle \overline{h'_\star} \rangle_{hu} = \begin{cases} a_{numAck} & \text{if } numAck > 0 \\ \langle \overline{h_\star} \rangle_{hu} & \text{otherwise} \end{cases}$$

This completes the definition of the behaviour of the circuit. We proceed with the description of an implementation.

As lemma 5.3 showed, the candidates can be found by looking first in the sorted elements of $\mathcal{R}^{>p}$ and then in the sorted elements of $\mathcal{R}^{\leq p}$. The implementation computes these sets by the equations

$$\begin{aligned} \text{rh}_\star &:= \text{req}_\star \wedge \text{h}_\star \\ \text{rl}_\star &:= \text{req}_\star \wedge \overline{\text{h}_\star} \end{aligned}$$

Since $\text{h}_\star = 1^{n-p}0^p$ we have:

$$\begin{aligned} \text{rh}_i = 1 &\iff i \in \mathcal{R}^{>p} \\ \text{rl}_i = 1 &\iff i \in \mathcal{R}^{\leq p} \end{aligned}$$

The search for the candidates for acknowledgement can now be reduced on finding the first k ones in the bus $(\text{rh}_\star, \text{rl}_\star)$. We use a FFk1hu_{2n} circuit, defined in section C.2. The FFk1hu_{2n} circuit returns a signal array $\text{o}_{\star,\star} \in \{0,1\}^{2n \cdot k}$. By corollary C.3 (p. 125) the columns $\text{o}_{\star,j}$ of $\text{m}_{\star,\star}$ being different from 0^{2n} deliver a prefix of the indices of the input bits equal to 1. This sequence is given in negated half-unary encoding: if $\langle \overline{\text{o}_{\star,j}} \rangle_{hu} = a$ the $(j+1)$ -th one is located at the input bit a . Due to the construction, positions indicated being greater or equal to n refer to indices in $\mathcal{R}^{>p}$ and have to be shifted back. Such a shift is indicated by $\text{o}_{n-1,k'} = 0$ since

$$\text{o}_{n-1,k'} = 0 \iff \langle \overline{\text{o}_{\star,k'}} \rangle_{hu} \geq n$$

The acknowledgement signals can therefore be computed the following way:³

$$\text{ack}_{k',\star} := (\text{o}_{n-1,k'} ? \text{o}_{n-1 \cdot 0, k'} : \text{o}_{2n-1 \cdot n, k'})$$

Finally, the register h_\star must be updated, if a candidate was acknowledged. Again, $\text{o}_{\star,\star}$ provides sufficient information to distinguish the different cases: the signal

$$\text{numAck}_\star := \text{o}_{2n-1,\star}$$

contains a half-unary encoding of the number of candidates found. We clock register h_\star in case that at least one candidate was found,

$$\text{hce} := \text{numAck}_0$$

Since the signal numAck_\star is a half-unary encoding of the number of candidates acknowledges, a pointer to the last acknowledged candidate can be found using an edge detector:

$$\text{lastAck}_\star := \text{markLastOne}(\text{numAck}_\star)$$

The update formula for h_\star can therefore be written the following way:

$$\text{h}'_\star = \bigvee \text{ack}_{k',\star} \wedge \text{lastAck}_i$$

Figure 5.1 shows an overview of the implementation just described. Much of the computation just outlined is hidden in the “update” box. For the initialization, h_\star is forced to 1^n .

Proof of Correctness. The correctness of the circuit follows from the construction of the circuit with help of lemma 5.3 and corollary C.3 as pointed out above.

³By a notational twist we also convert the columns $\text{o}_{\star,k'}$ to the rows $\text{ack}_{k',\star}$. The definition of $\text{ack}_{\star,\star}$ by rows is the more natural one.

5.2 Multicounter

A k -MCNT $_n$ is a counter that makes up to k steps in a single cycle. The step width is determined by an input $\text{inc}_\star \in \{0, 1\}^k$ containing a half-unary encoding of the number of steps. The counting register $\text{CNT}_\star \in \{0, 1\}^n$ satisfies the following equation:

$$\langle \text{CNT}'_\star \rangle_2 \equiv_{2^n} \begin{cases} 0 & \text{init} = 1 \\ \langle \text{CNT}_\star \rangle_2 + k' & \text{if } \langle \text{inc}_\star \rangle_{hu} = k' \end{cases}$$

Additionally, the multicounter returns the counters $\text{cnt}_{j,\star} \in \{0, 1\}^n$ for $0 \leq j \leq k$ with the property

$$\langle \text{cnt}_{j,\star} \rangle_2 = \langle \text{CNT}_\star \rangle_2 + j$$

An implementation of such a multicounter is shown in figure 5.2. An multiple incrementer (cf. section C.3) is used to compute the output counters $\text{cnt}_{j,\star}$ for $j \in \{0, \dots, k\}$. From these counters the correct one has to be selected for the update of CNT'_\star . Selection signals can be obtained by converting inc_\star into unary encoding using an edge detector:

$$\text{incU}_\star := \text{markLastOne}((\text{inc}_\star, 1))$$

The bus $\text{incU}_\star \in \{0, 1\}^{k+1}$ satisfies $\text{incU}_j = 1 \iff j = \langle \text{inc}_\star \rangle_{hu}$. Therefore, CNT'_\star can be computed with

$$\text{CNT}'_\star = \bigvee_{0 \leq j \leq k-1} \text{incU}_j \wedge \text{cnt}_{j,\star}$$

5.3 Multiported Queue

This section develops a multiported queue. A queue is a storage structure supporting read and write operations on a first-in-first-out (FIFO) basis. Multiporting allows multiple read / write accesses per cycle.

This section proceeds in two steps. First, an abstract model of a multiported queue is developed on the basis of finite state transducers. Second, an interface for the implementation is described which we use in two different implementations. The correctness of the implementations is verified by showing the equivalence to the abstract definition.

5.3.1 Definition of an Abstract Multiported Queue

Definition 5.4 A finite state transducer A is an 6-tuple $A = (I, Z, \delta, z_0, O, \epsilon)$; I is the input alphabet, Z is a set of states, $\delta : I \times Z \rightarrow Z$ is the state transition function, $z_0 \in Z$ is the initial state, O is the output alphabet, $\epsilon : I \times Z \rightarrow O$ is the output function.

Note that finite state transducers, as needed in this thesis, do not require a set of end states. Finite state transducers are only used to represent computations from one step to another.

Definition 5.5 (Abstract Multiported Queue) Let $k, l, n \in \mathbb{N}$, $k, l \leq n$, DOM a finite set. Define $K := \{0, \dots, k-1\}$, $L := \{0, \dots, l-1\}$, $N := \{0, \dots, n-1\}$. An

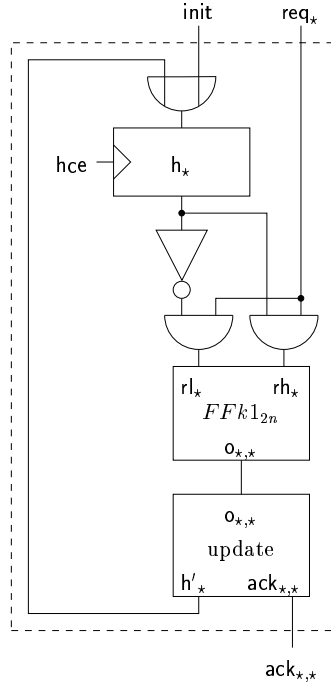


Figure 5.1: Implementation of a multiported round-robin selector

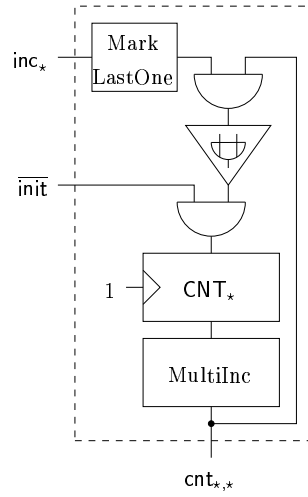


Figure 5.2: Definition of a multicounter




Figure 5.3: Read-write-operation on an abstract MPQ with $wack = w, rack = r$

abstract $k:l$ -multiported queue of size n over the domain DOM is a finite state transducer

$$\begin{aligned}
 A_{k:l-MPQ_n} &= (I, Z, \delta, z_0, O, \epsilon) \\
 \text{input alphabet } I &= K \times DOM^{\leq k} \times L \\
 \text{set of states } Z &= N \times DOM^{\leq n} \\
 \text{initial state } z_0 &= 0 \\
 \text{output alphabet } O &= K \times L \times DOM^{\leq l}
 \end{aligned}$$

The set of states Z models the storage space of the queue. The input alphabet I encodes the arrival of new objects and the output alphabet O encodes the reading-out of objects. The transition function $\delta : I \times Z \rightarrow Z$ and the output function $\epsilon : I \times Z \rightarrow O$ are defined as follows:

$$\begin{aligned}
 \delta(w, W_1, \dots, W_w, r, e, E_1, \dots, E_e) &= (e', E'_1, \dots, E'_{e'}) \\
 \epsilon(w, W_1, \dots, W_w, r, e, E_1, \dots, E_e) &= (wack, rack, R_1, \dots, R_{rack}) \\
 rack &:= \min\{r, e\} \\
 wack &:= \min\{n - e + rack, w\} \\
 e' &:= e + wack - rack \\
 (R_1, \dots, R_{rack}) &:= (E_1, \dots, E_{rack}) \\
 (E'_1, \dots, E'_{e'}) &:= (E_{rack+1}, \dots, E_e, W_1, \dots, W_{wack})
 \end{aligned}$$

It is easy to see, that this definition is a formalization of the concept of a multi-ported queue: the order of the queue entries (according to the insertion time) E_1, \dots, E_e is preserved. Entries are read head-first and written tail-first. Multiplicity is introduced by the read and write parameters. Figure 5.3 shows a graphical interpretation of these equations for the case that the queue is “nicely” filled (i.e. we have $wack = w, rack = r$). The first row shows the entries E_1, \dots, E_e being stored in the queue in cycle t . New data W_1, \dots, W_w arrives also in cycle t and r elements are requested for read-out. The second row shows the read-out entries R_1, \dots, R_r and the newly formed queue $E'_1, \dots, E'_{e'}$.

5.3.2 Interface

For the implementation consider the domain $DOM = \{0, 1\}^m$ for some $m \in \mathbb{N}$. Table 5.1 shows the definition of the implementation interface of a multiported queue. In addition to the interface for the read / write operations, we consider the information on how the data is stored in the queue also part of the interface. The queue data is stored in *slots* each of which is tagged by a valid bit. The entries E_1, \dots, E_e are to be found in the valid slots, by convention S_1, \dots, S_e . Note the

Name	Type	Meaning
$W_*.data_*$	In	the write data, busses for W_1, \dots, W_k
$W_*.req$	In	the write requests, half-unary encoding of r
$W_*.ack$	Out	the write acknowledgments, half-unary encoding of $rack$
$S_*.valid$	Out	half-unary encoding of e
$S_*.data_*$	Out	storage slots for the entries
$R_*.data_*$	Out	the read data, busses for R_1, \dots, R_l
$R_*.req$	In	the read requests, half-unary encoding of w
$R_*.ack$	Out	the read acknowledgments, half-unary encoding of $wack$

Table 5.1: Multi-ported queue interface

negation and reversion of the $E_*.valid$ bus can also be interpreted as a full bus, $full_* := \overline{E_{0..n-1}.valid}$. In this case, $full_i = 1$ iff i entries cannot be stored in queue (before read-out).

5.3.3 Register-Based Implementation

In the register-based implementation, the slots are directly modelled with flip-flops $S_i.data_j$ for $0 \leq i \leq n-1$ and $0 \leq j \leq m-1$. A valid register $valid_* \in \{0, 1\}^n$ indicates which slots hold valid entries. It encodes the fill state of the queue in half-unary encoding, i.e. $\langle valid_* \rangle_{hu}$ is the number of allocated queue slots.

Queue Control

This section develops the queue control signals in three different steps. The “Read Phase” describes how the read operations are acknowledged and performed. The “Slot Propagation” describes the adjustment of the queue slots after the read-out. The “Write Phase” describes how the write operations are acknowledged and how the data reaches the appropriate slots.

Read Phase. The reading of slots is simple. For the computation of the read acknowledgements signals we use the fact, that the minimum function used to compute $rack$ (cf. definition 5.5), is easily implemented in half-unary encoding with a slice of AND-gates:

$$R_*.ack := S_{l-1..0}.valid \wedge R_*.req$$

In preparation for the write phase, we assume that the read operations have been separately performed before the write operations. The fill state of the queue after read-out, $\langle validAfterRead_* \rangle_{hu}$ can be obtained by a half-unary subtraction operation:

$$\langle validAfterRead_* \rangle_{hu} = \langle valid_* \rangle_{hu} - \langle R_*.ack \rangle_{hu}$$

Slot Propagation. After reading out, the entries stored in the slots move down the queue to fill the read-out empty slots again. This movement is called slot propagation. Let $prop_{l'} = 1$ indicate that the propagation distance equals l' steps ($l' \in \{0, \dots, l\}$). Since $R_*.ack$ provides a half-unary encoding of the number of read-out entries, we can compute $prop_*$ with an edge detector:

$$prop_* := \text{markLastOne}(R_*.ack)$$

Write Phase. In the write phase, we first want to compute the write acknowledgement signals. According to property 3.6 (p. 35), of half-unary numbers, the reversed and negated version of $\text{validAfterRead}_\star$ encodes the number of free entries after read-out:

$$\langle \overline{\text{validAfterRead}_{0..n-1}} \rangle_{hu} = n - e + \text{rack}$$

Because $\text{wack} = \min\{n - e + \text{rack}, w\}$ by definition 5.5 we can compute the write acknowledgement signals by a slice of AND-gates:

$$\begin{aligned} W_\star.\text{ack} &= W_\star.\text{req} \wedge \overline{\text{validAfterRead}_{n-k..n-1}} \\ &= \overline{W_\star.\text{req} \wedge \text{validAfterRead}_{n-k..n-1}} \end{aligned}$$

The acknowledged write data must be appended to the valid entries in the queue after read-out. To control the distribution of the write busses, we define the signal array $W2S_{\star,\star}$. The condition $W2S_{k',n'} = 1$ indicates that write bus $W_{k'}$ shall be written to slot $S_{n'}$.

For $k' = 0$, the bus

$$\text{insPos}_\star := \text{markFirstZero}(\text{validAfterRead}_\star)$$

marks the insertion position, i.e. $W2S_{0,\star} = W_0.\text{req} \wedge \text{insPos}_\star$. For $k' > 0$ we obtain $W2S_{k',\star}$ by shifting insPos_\star :

$$W2S_{k',n'} = W_{k'}.\text{req} \wedge \text{insPos}_{n'-k'} \quad \text{for } 1 \leq n' - k' \leq n$$

This ensure that the written entries are appended “en bloc” after the valid queue entries.

Update of the Valid Register. The new valid register valid'_\star can be computed by a half-unary addition that satisfies:

$$\langle \text{valid}'_\star \rangle_{hu} = \langle \text{validAfterRead}_\star \rangle_{hu} + \langle W_\star.\text{ack} \rangle_{hu}$$

Figure 5.4 contains an overview of the calculations for the valid register. Additionally it initializes the valid_\star register if the signal init_\star equals 0.

Structure of the Slots

Figure 5.5 shows the definition of the slot $S_{n'}$. The register $S_{n'}.\text{data}_\star$ stores the data. Sources for this register are the k write busses $W_\star.\text{data}$, the own data $S_{n'}$ (which may be looped back for update purposes) and the data of the l above slots $S_{n'+1}, \dots, S_{n'+l}$ for slot propagation purposes. Drivers are used to select the data from the different sources; $W2S_{k',n'} = 1$ indicates, that write bus $W_{k'}$ shall write into slot $S_{n'}$. At most one bit of $W2S_{\star,n'}$ can be active by definition of $W2S_{\star,\star}$. The signal

$$S2S_{n'+l',n'} = \text{valid}_{n'+l'} \wedge \text{prop}_{l'}$$

for $l' \in \{0, \dots, l\}$ is used to select $S_{n'+l'}$ as a source. Again, at most one of $S2S_{n'+l',n'}$ can be active, since at most one of prop_\star is active by definition. Furthermore, the two different output enables busses are mutually exclusive:

$$\begin{aligned} \exists l' : S2S_{n'+l',n'} = 1 &\iff \text{prop}_{l'} = 1 \wedge \text{valid}_{n'+l'} = 1 \\ &\iff R_\star.\text{ack} = 0^{l-l'} 1^{l'} \wedge \text{valid}_{n'+l'} = 1 \\ &\iff \text{validAfterRead}_{n'} = 1 \\ &\iff W2S_{\star,n'} = 0 \end{aligned}$$

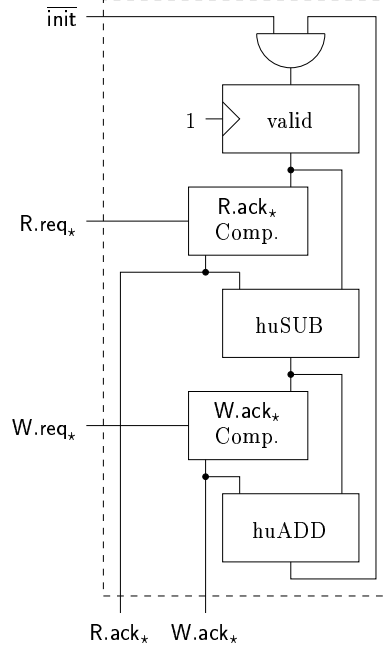



Figure 5.4: Definition of the valid register

Figure 5.5: Definition of the n' -th slot

The data register is only clocked, if it is valid in the next cycle, i.e. $S_{n'}.data.ce' = \text{valid}'_{n'}$.

Note that the construction of the slots allows the integration of an update circuit for the queue data, as shown in figure 5.5. Such a circuit can be used to modify entries after they have been written. However, for simplifying notation, we just consider the update circuit computing the identity function for the remaining part of this section.

Correctness

Our implementation behaves the same way as an abstract queue $A_{k,l-MPQ_n}$ on domain $DOM = \{0,1\}^m$. In an inductive proof, we show that the initialization configuration and the state transitions are equivalent subject to the interface defined in section 5.3.2.

For initialization we have $e = 0$ and no entries in the abstract machine. The implementation machine is in an equivalent state since

$$\begin{aligned} \langle \text{valid}_* \rangle_{hu} &= \langle 0^n \rangle_{hu} \\ &= 0 \\ &= e \end{aligned}$$

and the slots must not match any entries.

Now we show that both machines make equivalent transitions. Under the assumptions

$$\begin{aligned} W_*.req &= 0^{k-w} 1^w \\ W_{w-1..0}.data &= (W_w, \dots, W_1) \\ R_*.req &= 0^{l-r} 1^r \end{aligned} \tag{5.1}$$

the following equations hold:

$$\langle R_*.ack \rangle_{hu} = rack \tag{5.2}$$

$$R_{rack-1..0}.data = (E_{rack}, \dots, E_1) \tag{5.3}$$

$$\langle W_*.ack \rangle_{hu} = wack \tag{5.4}$$

$$\langle \text{valid}'_* \rangle_{hu} = e' \tag{5.5}$$

$$S_{e'-1..0}.data = (E'_{e'}, \dots, E'_1) \tag{5.6}$$

So, assume 5.1. First, we consider the read phase, for which we can directly derive by properties of half-unary encodings:

$$\begin{aligned} \langle R_*.ack \rangle_{hu} &= \langle R_*.req \wedge S_{l-1..0} \rangle_{hu} \\ &= \min \{r, e\} \\ &= rack \\ R_{rack-1..0}.data_* &= S_{rack-1..0}.data_* \\ &= (E_{rack}, \dots, E_1) \end{aligned}$$

Next, the equivalence of the write acknowledgement signals are shown. The following equations use properties of half-unary numbers and the correctness of the read acknowledgements already shown:

$$\begin{aligned} \langle W_*.ack \rangle_{hu} &= \langle W_*.req \wedge \overline{\text{validAfterRead}_{n-k..n-1}} \rangle_{hu} \\ &= \langle (0^{n-k}, W_*.req) \wedge \overline{\text{validAfterRead}_{0..n-k}} \rangle_{hu} \end{aligned}$$




Figure 5.6: Slot storage

$$\begin{aligned}
 &= \max \{w, n - \langle \text{validAfterRead}_{n-1..0} \rangle_{hu} \} \\
 &= \max \{w, n - \max \{0, e - rack\} \} \\
 &= \max \{w, n - e + rack\} \\
 &= wack
 \end{aligned}$$

The preservation of the queue slot structure remains to be shown. The updated valid register is computed by a half unary addition, so we have

$$\begin{aligned}
 \langle \text{valid}'_{*} \rangle_{hu} &= \langle \text{validAfterRead}_{*} \rangle_{hu} + \langle W_{*}.ack \rangle_{hu} \\
 &= e - rack + wack \\
 &= e'
 \end{aligned}$$

Each propagated slot moves down by a distance of $rack$. The written entries are stored in positions $e - rack + 1$ to $e - rack + wack$. So:

$$\begin{aligned}
 S_{e'-1..0}.data &= (W_{wack}, \dots, W_1, E_e, \dots, E_1) \\
 &= (E'_{e'}, \dots, E'_1)
 \end{aligned}$$

5.3.4 RAM-Based Implementation

The RAM-based implementation uses a regular multiported RAM accessed by head and tail pointers in a wrap-around fashion. This implementation is cheaper although slower than a register-based implementation. The implementation restricts for n being a power of 2. Sections on the structure of a RAM-based implementation follow.

Slot Storage

The slot storage is a multiported RAM with l write and k read ports. The read ports are addressed by the head pointers $head_{*,*}$ returning their data to the read busses $R_{*}.data_{*}$. The write ports are addressed by the tail pointers $tail_{*,*}$ and enabled by the write acknowledgement signals $W_{*}.ack$. The data is on the incoming write busses $W_{*}.data_{*}$. Figure 5.6 shows the slot storage environment.

Head and Tail Pointers

The queue is controlled a head and a tail pointer. The (master) head pointer $HEAD_{*}$ always points to the first (oldest) element in the slot storage. The (master) tail pointer $TAIL_{*}$ always points to the first free slot in the slot storage. The queue entries consist of the entries stored between head and tail. Addresses wrap around if they get larger than n .




Figure 5.7: Definition of the head and tail pointer

However, one head pointer and one tail pointer is not enough for multiple read and write operations. To support k read operations, we need k head pointers $\text{head}_{*,*}$ with

$$\langle \text{head}_{k',*} \rangle_2 = \langle \text{HEAD}_* \rangle_2 + k'$$

The k' -th head pointer can be used to read out the k' -th element. After reading out, the master head is incremented by $\langle \text{R}_*. \text{ack} \rangle_{hu}$.

The same holds for the tail pointers. To support l write operations, we need l tail pointers $\text{tail}_{*,*}$ with

$$\langle \text{tail}_{l',*} \rangle_2 = \langle \text{TAIL}_* \rangle_2 + l'$$

The l' -th tail pointer can be used to write the l' -th element. After writing, the master tail pointer is incremented by $\langle \text{W}_*. \text{ack} \rangle_{hu}$.

The multicounter, already implemented in section 5.2, provides this functionality for both the head and the tail pointer as shown in figure 5.7.

Valid Register

The valid register is implemented exactly the same way as for the register-based implementation. Refer to section 5.3.3 and figure 5.4 for details.

Correctness

We show correctness the same way as for the register-based implementation. First, we show that the initial configurations are equivalent and then we show inductively that equivalent transitions are performed. Equivalence is defined along with the interface of section 5.3.2. The definition of the slots of the RAM-based implementation is a slightly more complicated; the slots are located between head and tail of the queue. With $e = \langle \text{S}_*. \text{valid} \rangle_{hu}$ we satisfy

$$E_{e'} = S_{\text{head}+e'}$$

For the initialization we correctly have $e = 0$ and $\text{S}_*. \text{valid} = 0$. Then we show under the assumptions:

$$\begin{aligned} \text{W}_*. \text{req} &= 0^{k-w} 1^w \\ \text{W}_{w-1..0}. \text{data} &= (W_w, \dots, W_1) \\ \text{R}_*. \text{req} &= 0^{l-r} 1^r \end{aligned} \tag{5.7}$$

that the following equations hold:

$$\langle R_{\star}.ack \rangle_{hu} = rack \quad (5.8)$$

$$R_{rack-1..0}.data = (E_{rack}, \dots, E_1) \quad (5.9)$$

$$\langle W_{\star}.ack \rangle_{hu} = wack \quad (5.10)$$

$$\langle S_{\star}.valid' \rangle_{hu} = e' \quad (5.11)$$

$$S_{e'-1..0}.data = (E'_{e'}, \dots, E'_1) \quad (5.12)$$

We only show here that the data is correctly read and written to the queue. The other equations have already been proved in 5.3.3.

For read requests $r' \in \{0, \dots, rack - 1\}$, the r' -th head pointer is used. This pointer points to the r' element starting with and including the master head pointer $HEAD_{\star}$. Therefore by induction assumption:

$$R_{r'}.data = E_{r'+1}$$

Data from the write busses is written at the queue tail; $W_{w'}$ is written to the position indicated by the w' -th head pointer. With $T := \langle TAIL_{\star} \rangle_2$ we have

$$S_{T+w'}.data' = W_{w'+1}$$

Otherwise, the slot is not modified, so for $n' \notin \{T, \dots, T + wack\}$ we have

$$S'_{n'} = S_{n'}$$

Update Ports

The design of the RAM-based MPQ allows for the integration of *update ports*: these ports allow reading and writing data in the middle of the queue. To access an entry, its position is used as an address.

The update ports can be implemented by adding the required read and write ports to the slot storage RAM.

5.4 Reservation Station Queue

A reservation station queue is a specialized multiported queue maintaining an additional status bit **e4ro** (“eligible for read-out”) for each entry. If the bit is 1, an entry may be read; if it is 0 an entry cannot be read. This sort of behaviour is needed in collection of reservation stations: each reservation station becomes eligible for read-out, if it has gathered all its source operands.

We proceed the same way as for multiported queues: first we develop an abstract definition of a reservation station queue based on finite state transducers. A description of an implementation interface follows. Then, a register-based implementation is presented only.

5.4.1 Definition of an Abstract Multiported Queue

Definition 5.6 (Abstract Reservation Station Queue) *Let $k, l, n \in \mathbb{N}$, $k, l \leq n$, DOM a finite set. Define $K := \{0, \dots, n\}$, $L := \{0, \dots, l\}$, $N := \{0, \dots, n\}$. An abstract $k:l$ -reservation station queue of size n over the domain DOM is a finite state transducer*

$$\begin{aligned} A_{k:l-RSQ_n} &= (I, Z, \delta, z_0, O, \epsilon) \\ \text{input alphabet } I &= K \times DOM^{\leq k} \times L \end{aligned}$$

Name	Type	Meaning
$W_*.data_*$	In	the write data, busses for W_1, \dots, W_k
$W_*.req$	In	the write requests, half-unary encoding of r
$W_*.ack$	Out	the write acknowledgments, half-unary encoding of $rack$
$S_*.valid$	Out	half-unary encoding of e
$S_*.e4ro$	Out	the eligibility for read-out
$S_*.data_*$	Out	storage slots for the entries
$R_*.data_*$	Out	the read data, busses for R_1, \dots, R_l
$R_*.req$	In	the read requests, half-unary encoding of w
$R_*.ack$	Out	the read acknowledgments, half-unary encoding of $wack$

Table 5.2: Reservation station queue interface

$$\begin{aligned}
\text{set of states } Z &= N \times DOM^{\leq n} \\
\text{transition function } \delta &: I \times Z \longrightarrow Z \\
\text{initial state } z_0 &= 0 \\
\text{output alphabet } O &= K \times L \times DOM^{\leq l} \\
\text{output function } \epsilon &= I \times Z \longrightarrow O
\end{aligned}$$

Let $E4RO$ (“eligible for read-out”) be a finite predicate on DOM , $E4RO \subseteq DOM$. The transition function δ and the output function ϵ are defined as follows:

$$\begin{aligned}
\delta(w, W_1, \dots, W_w, r, e, E_1, \dots, E_e) &= (e', E'_1, \dots, E'_{e'}) \\
\epsilon(w, W_1, \dots, W_w, r, e, E_1, \dots, E_e) &= (wack, rack, R_1, \dots, R_{rack})
\end{aligned}$$

Define $e4ro := |\{i \in \{1, \dots, e\} \mid E_i \in E4RO\}|$. Then

$$\begin{aligned}
rack &:= \min \{rack, e4ro\} \\
wack &:= \min \{n - e + rack, w\} \\
e' &:= e + wack - rack
\end{aligned}$$

Let $(e4ro_j)$ be the prefix of length $rack$ of the sorted sequence of the indices of the elements eligible for read out. Let $(\overline{e4ro}_k)$ be the sorted sequence of the remaining indices. Then, with an update function upd for the entries, the queue operation looks as follows:

$$\begin{aligned}
(R_1, \dots, R_{rack}) &:= (E_{e4ro_1}, \dots, E_{e4ro_{rack}}) \\
(E'_1, \dots, E'_{e'}) &:= \left(upd(E_{\overline{e4ro}_1}), \dots, upd(E_{\overline{e4ro}_{e-rack}}), W_1, \dots, W_{wack} \right)
\end{aligned}$$

5.4.2 Description of the Interface and Equivalence Criterion

For the implementation of the abstract multiported queue we consider a domain $DOM = \{0, 1\}^m$ for some $m \in \mathbb{N}$. Table 5.2 shows the definition of the implementation interface of a multiported queue.

5.4.3 Register-Based Implementation

In the register-based implementation, the slots are directly modelled with flip-flops $S_*.data_*$. A valid register $valid_* \in \{0, 1\}^n$ indicates which slots hold valid entries. It encodes the fill state of the queue in half-unary encoding, i.e. $\langle valid_* \rangle_{hu}$ is the number of allocated queue slots.

Queue Control

This section develops the queue control signals in three steps. The procedure is according to the multiported queue case. The “Read Phase” describes how queue read operation are acknowledged and performed. The “Entry Propagation” describes the adjustment of the queue entries after the read-out. The “Write Phase” describes how queue write operations are acknowledged and how the data reaches the appropriate slots.

Read. Candidates for read-out must be searched in all valid and eligible entries:

$$S_{\star}.cand = S_{\star}.valid \wedge S_{\star}.e4ro$$

The definition of the abstract reservation station queue requires us to find the first l ones in this bit string. A find-first- l -ones half-unary circuit executes this operation. Let $S_{\star}.cand \in \{0,1\}^n$ be the input of this circuit, $mAux_{\star,\star} \in \{0,1\}^{n \cdot l}$ the output. We have

$$\langle mAux_{i,\star} \rangle_{hu} = \min \{l, \text{ones}(S_{i..0}.cand)\}$$

by definition. As we are interested to determine the number of read-out entries, we compute a half-unary minimum operation

$$m_{i,\star} := R_{\star}.req \wedge mAux_{i,\star}$$

and obtain

$$\begin{aligned} \langle m_{i,\star} \rangle_{hu} &= \min \{ \langle R_{\star}.req \rangle_{hu}, l, \text{ones}(S_{i..0}.cand) \} \\ &= \min \{ \langle R_{\star}.req \rangle_{hu}, \text{ones}(S_{i..0}.cand) \} \end{aligned}$$

From this signal array, the read acknowledgement signals can be directly read out from the last line:

$$R_{\star}.ack := m_{n-1,\star}$$

To compose the read busses we further need to compute output enable signals $S2R_{i,j}$ indicating that slot i shall be mapped to read bus j . Again, the signal array $m_{\star,\star}$ provides this information. By corollary C.3 (p. 125) the columns $m_{\star,j}$ of $m_{\star,\star}$ being different from 0ⁿ deliver a $\langle R_{\star}.ack \rangle_{hu}$ -prefix of the sorted sequence of the indices of the input bits equal to 1. This sequence is given in negated half-unary encoding; to obtain the output signals, we merely compute the edge in them:

$$S2R_{\star,j} := \text{markLastOne}(m_{\star,j})$$

The rest of the read phase works the same as in the multiported queue. In preparation of the write phase, we assume that the read operations have been separately performed before the write operations. The signal bus $validAfterRead_{\star}$ indicating the valid slots after the completion of the read phase is obtained by a half-unary subtraction operation:

$$\langle validAfterRead_{\star} \rangle_{hu} = \langle valid_{\star} \rangle_{hu} - \langle R_{\star}.ack \rangle_{hu}$$

Slot Propagation. The entry propagation in a reservation station queue is more complicated than in a multiported queue. Each entry may be propagated by a different distance since reading of queue entries can occur at different places of the queue and is not restricted to the head.

Let $S_i.\text{prop}_j$ indicate that the propagation distance of entry i is j . The slot propagation is characterized by two rules. A read-out slot is not propagated at all. Any other slot is propagated by the number of entries that have been read out before it.

Define

$$S_i.\text{prop}_\star := m_{i-1,\star} \wedge \text{markLastOne}(m_{i,\star})$$

Then, for any read-out slot S_i , we have $\langle m_{i-1,\star} \rangle_{hu} < \langle m_{i,\star} \rangle_{hu}$ and therefore in $S_i.\text{prop}_\star = 0^l$. Otherwise,

$$\begin{aligned} \langle S_i.\text{prop}_\star \rangle_{hu} &= \text{ones}(S_{i..0}.\text{cand}) \\ &= \text{ones}(S_{i-1..0}.\text{cand}) \end{aligned}$$

Write. The write phase works the same as in the multiported queue (cf. section 5.3.3). In the write phase, we first want to compute the write acknowledgement signals. According to property 3.6 (p. 35) of half-unary encodings, the reversed and negated version of $\text{validAfterRead}_\star$ encodes the number of free entries after read-out:

$$\overline{\langle \text{validAfterRead}_{0..n-1} \rangle_{hu}} = n - e + \text{rack}$$

Because $wack = \min\{n - e + \text{rack}, w\}$ by definition 5.5, we can compute the write acknowledgement signals by a slice of AND-gates:

$$\begin{aligned} W_\star.\text{ack} &= W_\star.\text{req} \wedge \overline{\text{validAfterRead}_{n-k..n-1}} \\ &= \overline{W_\star.\text{req} \wedge \text{validAfterRead}_{n-k..n-1}} \end{aligned}$$

The acknowledged write data must be appended to the valid entries in the queue after read-out. To control the distribution of the write busses, we define the signal array $W2S_{\star,\star}$. The condition $W2S_{k',n'} = 1$ indicates that write bus $W_{k'}$ shall be written to slot $S_{n'}$.

For $k' = 0$, the bus

$$\text{insPos}_\star := \text{markLastOne}(\text{validAfterRead}_\star)$$

marks the insertion position, i.e. $W2S_{0,\star} = W_0.\text{req} \wedge \text{insPos}_\star$. For $k' > 0$ we obtain $W2S_{k',\star}$ by shifting insPos_\star :

$$W2S_{k',n'} = W_i.\text{req} \wedge \text{insPos}_{n'-k'} \text{ for } 1 \leq n' - k' \leq n$$

This ensures that the written entries are appended “en bloc” after the valid queue entries.

Update of the Valid Register. The new valid register valid'_\star can be computed by a half-unary addition that satisfies:

$$\langle \text{valid}'_\star \rangle_{hu} = \langle \text{validAfterRead}_\star \rangle_{hu} + \langle W_\star.\text{ack} \rangle_{hu}$$

Entry

Figure 5.8 shows the definition of the slot $S_{n'}$. The register $S_{n'}.\text{data}_\star$ stores the data. Sources for this register are the k write busses $W_\star.\text{data}$, the own data $S_{n'}$ looped-back for update purposes and the data of the l above slots $S_{n'+1}, \dots, S_{n'+l}$ for slot propagation purposes. Drivers are used to select the data from the different




Figure 5.8: Definition of a reservation station queue slot

sources; $W2S_{k',n'} = 1$ indicates that write bus $W_{k'}$ shall write into slot $S_{n'}$. At most one bit of $W2S_{*,n'}$ can be active by definition of $W2S_{*,*}$. The signal

$$S2S_{n'+l',n'} := \text{valid}_{n'+l'} \wedge S_{n'+l'}.prop_{l'}$$

for $l' \in \{0, \dots, l\}$ is used to select $S_{n'+l'}$ as a source. Again, at most one of $S2S_{n'+l'}.n'$ can be active. The data register is only clocked, if it is valid in the next cycle, i.e. $S_{n'}.data.ce := \text{valid}'_{n'}$.

Chapter 6

Perspective

Several aspect of superscalar processor design have not been examined in-depth in this thesis. They are left for further research.

- Modern microprocessors execute multiple control flows with several outstanding branches in parallel. Such execution can be implemented by introduction of a *control flow tag*, which associates each instruction in execution with a specific control flow followed. The reorder buffer will only retire the instructions of the control flow that eventually is know to execute, results for other control flows will be dropped. Duplication of the producer tables is necessary for an individual register renaming and forwarding in each control flow.

Multiple control flow execution therefore depends on major changes of the machine design; an analysis would be of interest.

- More elaborate branch predictors and fetch mechanism still remain to be examined. [Yeh93] states the importance of such mechanism with the growing instruction-level parallelism of the processor since mispredictions incur a great penalty in performance. Therefore, current processor designs spend more cost for sophisticated control flow predictors.
- Branch problem “solved”, the memory operations are likely to limit the machine’s performance. In our design, stores are executed in order and loads execute on completed address computation. More elaborate schemes, as developed for example in [AMS97], involve the speculation of memory access addresses by maintaining appropriate history information. Schemes replacing in-order store by something more efficient are also thinkable, although preciseness for such schemes regarding interrupt and roll-back is a most delicate matter.
- Our design provided a superscalar framework for the whole instruction set. In contrast to this, today’s microprocessor often provide different levels of superscalarity for fixed-point and for floating point instructions. Additionally, the grouping of reservation stations in a centralized reservation station pool / dispatch stack results in a better utilization of reservation stations on a cost overhead ([Joh91, WS84, AKT86]).

The effect of such choices on cost and performance remains to be examined.

- Simulations are missing to measure the exact impact of the parameters of our design. The conservative CPI rate estimate made in chapter 4 could be could be corrected, improving the quality of this design.

- The hardware model used neglects the effect of wiring and fanout. A study of this design in a model dealing with layout, like [PS98], could adequately cope with the large wiring overhead present in our design.

Appendix A

DLX Instruction Set Architecture

This instruction set is taken from [MP95, MP00] with minimal modifications.




Figure A.1: Instruction formats of the DLX

$l_{31..26}$	Mnemonic	Effect
Control Flow Operation		
000010	j	$\langle PC_* \rangle_2 =_{2^{32}} \langle PC_* \rangle_2 + 4 + [\text{imm26}_*]$
000011	jal	$\langle R31_* \rangle_2 =_{2^{32}} \langle PC_* \rangle_2 + 4, \quad PC_* = \text{imm26}_*$
Exception Control		
111110	trap	$EPC_* = PC_*, PC_* = SISR_*$ $ESR_* = SR_*, ECA_* = MCA_*$ $SR_* = 0^{32}, EDATA_* = \text{sext}_{32}(\text{imm26}_*)$ clear CA but catch new interrupt events
111111	rfe	$SR_* = ESR_*; PC_* = EPC_*$

Table A.1: J-type instructions

$ _{31..26}$	Mnemonic	Effecto
Load / Store, $\text{mem}_* = M[\langle \text{RS1}_* \rangle_2 + [\text{imm16}_*]]$		
1000000	lb	$\text{RD}_* = \text{sext}_{32}(\text{mem}_{7..0})$
100001	lh	$\text{RD}_* = \text{sext}_{32}(\text{mem}_{15..0})$
100011	lw	$\text{RD}_* = \text{mem}_{31..0}$
100100	lbu	$\text{RD}_* = (0^{24}, \text{mem}_{7..0})$
100101	lhu	$\text{RD}_* = (0^{16}, \text{mem}_{15..0})$
101000	sb	$\text{mem}_{7..0} = \text{RD}_{7..0}$
101001	sh	$\text{mem}_{15..0} = \text{RD}_{15..0}$
101011	sw	$\text{mem}_{31..0} = \text{RD}_*$
Arithmetic / Logical Operation		
001000	addi	$\langle \text{RD}_* \rangle_2 =_{2^{32}} [\text{RS1}_*] + [\text{imm16}_*]$
001001	addiu	$\langle \text{RD}_* \rangle_2 =_{2^{32}} [\text{RS1}_*] + [\text{imm16}_*]$ (no overflow)
001010	subi	$\langle \text{RD}_* \rangle_2 =_{2^{32}} [\text{RS1}_*] + [\text{imm16}_*]$
001011	subiu	$\langle \text{RD}_* \rangle_2 =_{2^{32}} [\text{RS1}_*] + [\text{imm16}_*]$ (no overflow)
001100	andi	$\text{RD}_* = \text{RS1}_* \wedge \text{sext}_{32}(\text{imm16}_*)$
001101	ori	$\text{RD}_* = \text{RS1}_* \vee \text{sext}_{32}(\text{imm16}_*)$
001110	xori	$\text{RD}_* = \text{RS1}_* \oplus \text{sext}_{32}(\text{imm16}_*)$
001111	lhgi	$\text{RD}_* = (\text{imm16}_*, 0^{16})$
Test Set Operation ^a		
011000	clri	$\text{RD}_* = 0^{32}$
011001	sgri	$\text{RD}_* = ([\text{RS1}_*] > [\text{imm16}_*] ? 0^{31}1 : 0^{32})$
011010	seqi	$\text{RD}_* = ([\text{RS1}_*] = [\text{imm16}_*] ? 0^{31}1 : 0^{32})$
011011	sgei	$\text{RD}_* = ([\text{RS1}_*] \geq [\text{imm16}_*] ? 0^{31}1 : 0^{32})$
011100	slsi	$\text{RD}_* = ([\text{RS1}_*] < [\text{imm16}_*] ? 0^{31}1 : 0^{32})$
011101	snei	$\text{RD}_* = ([\text{RS1}_*] \neq [\text{imm16}_*] ? 0^{31}1 : 0^{32})$
011110	slei	$\text{RD}_* = ([\text{RS1}_*] \leq [\text{imm16}_*] ? 0^{31}1 : 0^{32})$
011111	seti	$\text{RD}_* = 0^{31}1$
Control Flow Operation		
000100	beqz	$\langle \text{PC}_* \rangle_2 =_{2^{32}} \langle \text{PC}_* \rangle_2 + 4 + (\text{RS}_* = 0^{32} ? [\text{imm16}_*] : 0)$
000101	bnez	$\langle \text{PC}_* \rangle_2 =_{2^{32}} \langle \text{PC}_* \rangle_2 + 4 + (\text{RS}_* \neq 0^{32} ? [\text{imm16}_*] : 0)$
010110	jr	$\text{PC}_* = \text{RS1}_*$
010111	jalr	$\langle \text{R31}_* \rangle_2 =_{2^{32}} \langle \text{PC}_* \rangle_2 + 4, \quad \text{PC}_* = \text{RS1}_*$

^aThe six relations defined are redundant for compiler-generated code in most situations: Compilers can replace relations “<”, “≠” and “≤” by their inverted operations in arithmetic expressions. The redundant operations could be replaced by unsigned comparisons which miss in the ISA presented. We will not follow this approach, however, to maintain binary compatibility to [MP95, MP00] and because the underlying ISA is not of primary interest of the thesis.

Table A.2: I-type instructions

$l_{31..26}$	$l_{5..0}$	Mnemonic	Effect
Shift Operation			
000000	000000	slli	$RD_* = RS1_* \ll SA_*$
000000	000001	slai	$RD_* = RS1_* \ll SA_*$ (arithmetic shift)
000000	000010	srli	$RD_* = RS1_* \gg SA_*$
000000	000011	srai	$RD_* = RS1_* \gg SA_*$ (arithmetic shift)
000000	000100	slli	$RD_* = RS1_* \ll RS_{4..0}$
000000	000101	slai	$RD_* = RS1_* \ll RS_{4..0}$ (arithmetic shift)
000000	000110	srli	$RD_* = RS1_* \gg RS_{4..0}$
000000	000111	srai	$RD_* = RS1_* \gg RS_{4..0}$ (arithmetic shift)
Data Transfer Operation			
000000	010000	movs2i	$RD_* = SA_*$
000000	010001	movi2s	$SA_* = RD_*$
Arithmetic / Logical Operation			
000000	100000	add	$\langle RD_* \rangle_2 =_{2^{32}} [RS1_*] + [RS2_*]$
000000	100001	addu	$\langle RD_* \rangle_2 =_{2^{32}} [RS1_*] + [RS2_*]$ (no overflow)
000000	100010	sub	$\langle RD_* \rangle_2 =_{2^{32}} [RS1_*] - [RS2_*]$
000000	100011	subu	$\langle RD_* \rangle_2 =_{2^{32}} [RS1_*] - [RS2_*]$ (no overflow)
000000	001100	and	$RD_* = RS1_* \wedge RS2_*$
000000	001101	or	$RD_* = RS1_* \vee RS2_*$
000000	001110	xor	$RD_* = RS1_* \oplus RS2_*$
000000	001110	lhg	$RD_* = (RS2_*, 0^{16})$
Test Set Operation			
000000	101000	clri	$RD_* = 0^{32}$
000000	101001	sgri	$RD_* = ([RS1_*] > [RS2_*]) ? 0^{31}1 : 0^{32}$
000000	101010	seqi	$RD_* = ([RS1_*] = [RS2_*]) ? 0^{31}1 : 0^{32}$
000000	101011	sgei	$RD_* = ([RS1_*] \geq [RS2_*]) ? 0^{31}1 : 0^{32}$
000000	101100	slsi	$RD_* = ([RS1_*] < [RS2_*]) ? 0^{31}1 : 0^{32}$
000000	101101	snei	$RD_* = ([RS1_*] \neq [RS2_*]) ? 0^{31}1 : 0^{32}$
000000	101110	slei	$RD_* = ([RS1_*] \leq [RS2_*]) ? 0^{31}1 : 0^{32}$
000000	101111	seti	$RD_* = 0^{31}1$

Table A.3: R-type instructions

$l_{31..26}$	Mnemonic	Effect
Load / Store, $\text{mem}_{63..0} = M[\langle RS1_* \rangle_2 + [imm16_*]]$		
110001	load.s	$FD_{31..0} = \text{mem}_{31..0}$
110101	load.d	$FD_{63..0} = \text{mem}_{63..0}$
111001	store.s	$\text{mem}_{31..0} = FD_{31..0}$
111101	store.d	$\text{mem}_{63..0} = FD_{63..0}$
Control Flow Operation		
000110	fbeqz	$\langle PC_* \rangle_2 =_{2^{32}} \langle PC_* \rangle_2 + 4 + (FCC = 0 ? [imm16_*] : 0)$
000111	fbnez	$\langle PC_* \rangle_2 \neq_{2^{32}} \langle PC_* \rangle_2 + 4 + (FCC = 1 ? [imm16_*] : 0)$

Table A.4: FI-type instructions

$l_{31..26}$	$l_{5..0}$	$l_{8..6}$	Mnemonic	Effect
Arithmetic / Compare Operation, $\text{fmt} := l_{8..6}$				
010001	000000		<code>fadd.s/d</code>	$\text{FD}_* = \text{ieeeAdd}(\text{FS1}_*, \text{FS2}_*, \text{fmt}_*)$
010001	000001		<code>fsub.s/d</code>	$\text{FD}_* = \text{ieeeSub}(\text{FS1}_*, \text{FS2}_*, \text{fmt}_*)$
010001	000010		<code>fmul.s/d</code>	$\text{FD}_* = \text{ieeeMul}(\text{FS1}_*, \text{FS2}_*, \text{fmt}_*)$
010001	000011		<code>fdiv.s/d</code>	$\text{FD}_* = \text{ieeeDiv}(\text{FS1}_*, \text{FS2}_*, \text{fmt}_*)$
010001	000100		<code>fneg.s/d</code>	$\text{FD}_* = \text{ieeeNeg}(\text{FS1}_*, \text{fmt}_*)$
010001	000100		<code>fabs.s/d</code>	$\text{FD}_* = \text{ieeeAbs}(\text{FS1}_*, \text{fmt}_*)$
010001	000100		<code>fsqt.s/d</code>	$\text{FD}_* = \text{ieeeSqrt}(\text{FS1}_*, \text{fmt}_*)$
010001	000100		<code>frem.s/d</code>	$\text{FD}_* = \text{ieeeRem}(\text{FS1}_*, \text{FS2}_*, \text{fmt}_*)$
010001	<code>11c_{3..0}</code>		<code>fc.cond.s/d</code>	$\text{FCC} = \text{ieeeCond}(\text{FS1}_*, \text{FS2}_*, \text{c}_*, \text{fmt}_*)$
Data Transfer Operation				
010001	001000	000	<code>fmov.s</code>	$\text{FD}_{31..0} = \text{FS1}_{31..0}$
010001	001000	001	<code>fmov.d</code>	$\text{FD}_{63..0} = \text{FS1}_{63..0}$
010001	001001		<code>mf2i</code>	$\text{RS}_* = \text{FS1}_{31..0}$
010001	001010		<code>mi2f</code>	$\text{FD}_{31..0} = \text{RS}_*$
Conversion				
010001	100000	001	<code>cvt.s.d</code>	$\text{FD}_* = \text{ieeeConv}(\text{FS1}_*, s, d)$
010001	100000	100	<code>cvt.s.i</code>	$\text{FD}_* = \text{ieeeConv}(\text{FS1}_*, s, i)$
010001	100001	000	<code>cvt.d.s</code>	$\text{FD}_* = \text{ieeeConv}(\text{FS1}_*, d, s)$
010001	100001	100	<code>cvt.d.d</code>	$\text{FD}_* = \text{ieeeConv}(\text{FS1}_*, d, i)$
010001	100100	000	<code>cvt.i.s</code>	$\text{FD}_* = \text{ieeeConv}(\text{FS1}_*, i, s)$
010001	100100	001	<code>cvt.i.d</code>	$\text{FD}_* = \text{ieeeConv}(\text{FS1}_*, i, d)$

Table A.5: FR-type instructions

Appendix B

Sample Branch Predictor Unit

B.1 Branch Predictor

We only implement a simple branch predictor in our design. The design of more a complex branch predictor lies beyond the scope of this thesis. Our branch predictor is a branch target buffer design using a 2-bit saturating counter scheme. The description and analysis such schemes can be found in [LS84].¹

B.1.1 Implementation

Each branch, identified by its address, is associated with a status information $(c_1, c_0) \in \{0, 1\}^2$. This information is the state of an automaton called 2-bit saturating counter. Figure B.1 shows this automaton. The automaton is used in two ways. First, branches are predicted based on the state (c_1, c_0) . A branch is predicted taken (predT) if $c_1 = 1$; it is predicted fall-through (predFT) if $c_1 = 0$. After the prediction of a branch, the state is updated according to the real outcome of that branch. A taken result increments the state number $\langle c_* \rangle_2$ up to its maximum while a fall-through result decrements the state number $\langle c_* \rangle_2$, but not below 0. This way, the counter “saturates” for many taken or for many fall-through results. A simple proof shows that the next state (c'_1, c'_0) can be computed out of (c_1, c_0) and T (for taken result) with the following equations:

$$\begin{aligned} c'_0 &:= c_1 \wedge \overline{c_0} \vee \overline{(c_1 \oplus c_0)} \wedge T \\ c'_1 &:= c_1 \wedge \overline{T} \vee c_0 \wedge T \end{aligned}$$

Figure B.2 shows the data paths for a 2-bit saturating counter branch predictor. The predictor is interfaced by 6 signals. Requests for predictions are made by setting $R = 1$ and ID to the ID of the branch. We identify branches in a direct-mapped approach by part of its PC. The predictor outputs the prediction result with the signal $PRED$; $PRED = 1$ indicates a taken branch. If a prediction is actually used, the signal DO must be activated.

The remaining part of the interface deals with the verification of predictions. Signal WB indicates that a branch has been verified. In this case MP is set to 1 if a misprediction has been detected.

¹Note that the DLX only implements branches with *static targets*, i.e. the target is encoded as a constant offset to the PC of the branch instructions. The original branch target buffer design, presented by Lee and Smith, also treated *dynamic targets* by storing an additional target item, giving the buffer its name. Dropping this target item simplifies the description of the scheme.




Figure B.1: State automaton for the 2-bit saturating counter

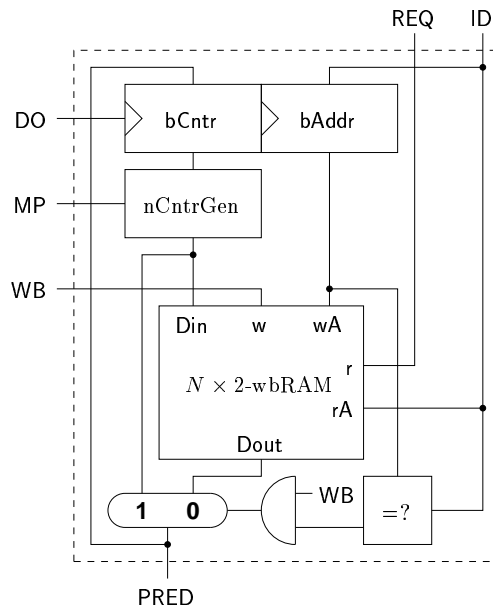


Figure B.2: Data paths for a 2-bit saturating counter branch predictor

The predictor features a central $N \times 2$ -wbRAM (see below) storing N 2-bit saturating counters. If a prediction is requested, this RAM is read out at the address ID_* , returning the counter (c_1, c_0) of the requested branch. From this counter, we can directly take the upper bit c_1 to obtain the prediction **PRED**. In case this prediction is actually used, i.e. $DO = 1$, the address ID_* and the counter (c_1, c_0) are clocked into two buffer registers $bAddr_*$ and $bCntr_*$. There it waits for the prediction verification.

If $WB = 1$, the verification result **MP** is available. The new counter (c'_1, c'_0) can be computed by the equations above. Note that $T = c_1 \oplus MP$. A write access to address $bAddr_*$ writes back the counter. The writing counter is forwarded for reading, if it is requested in the same round: so if $ID_* = bAddr_*$ and $WB = 1$ then (c'_1, c'_0) is returned.

A close observation of our processor reveals that the underlying RAM of the branch predictor need not have two “real” ports for reading and writing. A cycle after prediction, our processor will not request for another prediction, because it first has to resolve the outstanding speculation. This means, that between two possible write-backs due to prediction verification there is always a cycle where no prediction is requested, i.e. no read access occurs. So-called write-buffered RAM, implemented in section B.2, exploits this property and is used to build a fast counter RAM for the predictor.

B.1.2 Integration in the processor

The interface of the BPU developed in the previous section are defined as follows:

$$\begin{aligned}
 BPU.ID_* &:= cfi.opc_{31..2} \\
 BPU.REQ &:= cfi.valid \\
 BPU.DO &:= cfi.doPred \\
 BPU.MP &:= bcu.mp \\
 BPU.WB &:= bcu.valid \wedge \overline{bcu.jump} \\
 bpu.pred &:= cfi.t.imm16
 \end{aligned}$$

Note that `allowCFI` will be zero the cycle after a prediction has been made.

B.2 Write-Buffered RAM

This section describes the implementation of a special 2-ported RAM that requires a cycle of no read requests between two write requests. Such RAM can be built with conventional single-ported RAM and an additional buffer for write operations.

Figures B.3 and B.4 shows the implementation of a $N \times m$ -write-buffered RAM featuring the central $N \times m$ -RAM. We continue describing the write and the read operation.

B.2.1 Write

Let $n := \lceil \log_2 N \rceil$. Written data is always stored in a buffer register first; $bAddr_* \in \{0, 1\}^n$ buffers the write address, $bDin_* \in \{0, 1\}^m$ buffers the incoming data. An additional register `valid` $\in \{0, 1\}$ equals one, if the buffer contains data to be written.

A full buffer is written to the RAM, if no read request has to be fulfilled:

$$wb := valid \wedge \bar{r}$$

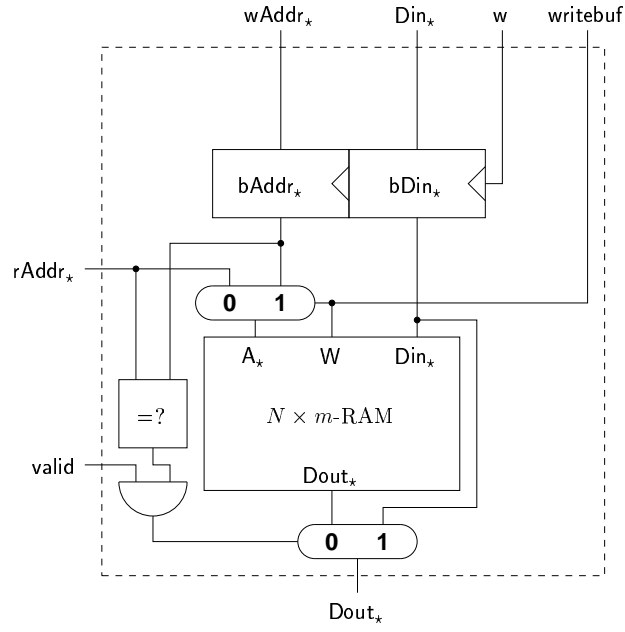


Figure B.3: Implementation of write-buffered RAM




Figure B.4: Implementation of the valid register

In this case the write signal is set and the buffered write address is selected as the RAM access address. The buffer valid register `bValid` is set to 1 iff it was valid and has not yet been written or if new write data has arrived in the same cycle:

$$\text{valid}' := (\text{bValid} \wedge \overline{\text{wb}}) \vee \text{w}$$

On initialization `valid` is reset.

B.2.2 Read

As seen above, no write operation will be executed, if a read is requested by `r = 1`. The address multiplexor will forward the read address to the address input of the RAM. Additionally, a forwarding logic is implemented for the data stored in the write buffer. A compare circuit “=” is used to determine matching read and buffer address. If this circuit signals matching addresses *and* the buffer is valid, read data is not taken from the RAM but forwarded from the buffer data register `bDin`.

Appendix C

Auxiliary Circuits

This appendix describes three auxiliary circuits used in the implementation of our machine.

C.1 Find-First-One Half-Unary

C.1.1 Definition and Construction

This section implements a find-first-one half-unary circuit FF1hu_n frequently used in this thesis. The FF1hu_n receives an input bus $\mathbf{a}_\star \in \{0,1\}^n$ and produces the output bus $\mathbf{o}_\star \in \{0,1\}^n$ with the following property:

$$\mathbf{o}_i = 1 \iff \exists j \in \{0, \dots, i\} : \mathbf{a}_j = 1$$

Equivalently, the negated output is the number of leading zeroes at the beginning of \mathbf{a}_\star in half-unary encoding:

$$\langle \overline{\mathbf{o}_\star} \rangle_2 = k \iff \mathbf{a}_{k..0} = 10^k$$

For $n = 1$, FF1hu_n is defined by $\mathbf{o}_0 := \mathbf{a}_0$. For $n > 1$, figure C.1 shows the recursive definition. It has delay of $\mathcal{O}(\log n)$ and cost of $\mathcal{O}(n \log n)$. Alternatively, a find-first-one half-unary circuit can be implemented with a parallel-prefix OR circuit (cf. [Lei99]).

C.1.2 Correctness.

The OR gate used in the construction can also be seen as the first bit of a half-unary adder for two 1-bit encodings. This way, the correctness proof of section C.2 for a generalized version of the circuit applies.

C.2 Find-First- k -Ones

C.2.1 Definition

Let the function $\text{ones}(\cdot)$ sum up the ones of a binary argument, i.e.

$$\text{ones}(\mathbf{x}_{m-1..0}) := \sum_{i < m} \mathbf{x}_i$$

Let $k, n \in \mathbb{N}$, $\mathbf{a}_\star \in \{0,1\}^n$. Define

$$O_j := \min \{k, \text{ones}(\mathbf{a}_{j..0})\}$$

On input \mathbf{a}_\star a n -bit find-first- k -ones half-unary circuit FFk1hu_n computes the output array $\mathbf{o}_{\star,\star} \in \{0,1\}^{n \cdot k}$ such that

$$\langle \mathbf{o}_{j,\star} \rangle_{hu} = O_j$$

C.2.2 Construction

For $k \geq n = 1$ the FFk1hu_n is defined by

$$\begin{aligned} \mathbf{o}_{1,1} &:= \mathbf{a}_1 \\ \mathbf{o}_{1,i} &:= 0 \end{aligned}$$

Now assume the construction of FFk1hu_n . We construct FFk1hu_{2n} using two FFk1hu_n and n half-unary adders (cf. equation 3.14) in the following way. One FFk1hu_n receives

$$\mathbf{a}l_\star := \mathbf{a}_{n-1..0}$$

and computes $\mathbf{o}l_{\star,\star} \in \{0,1\}^{n \cdot k}$. The other receives

$$\mathbf{a}h_\star := \mathbf{a}_{2n..n}$$

and computes $\mathbf{o}h_{\star,\star} \in \{0,1\}^{n \cdot k}$. The outputs of the FFk1hu_{2n} are computed such:

$$\mathbf{o}_{j,\star} := \begin{cases} \mathbf{o}l_{j,\star} & \text{for } j \in \{0, \dots, n-1\} \\ \text{huAdd}(\mathbf{o}h_{j-n,\star}, \mathbf{o}l_{n-1,\star}) & \text{for } j \in \{n, \dots, 2n-1\} \end{cases}$$

Only the k lower bits of the add operations are used; this puts an upper bound to the values $\langle \mathbf{o}_{j,\star} \rangle_{hu}$ computed as in the definition of O_j .

C.2.3 Correctness

In an inductive proof we assume the correctness of FFk1hu_n to show the correctness of FFk1hu_{2n} . According to the construction of FFk1hu_{2n} we distinguish two cases.

- $j \in \{0, \dots, n-1\}$. We have:

$$\begin{aligned} \langle \mathbf{o}_{j,\star} \rangle_{hu} &= \langle \mathbf{o}l_{j,\star} \rangle_{hu} \\ &= O_j \end{aligned}$$

- $j \in \{n, \dots, 2n-1\}$. We have:

$$\begin{aligned} \langle \mathbf{o}_{j,\star} \rangle_{hu} &= \min \{k, \langle \text{huAdd}(\mathbf{o}h_{j-n,\star}, \mathbf{o}l_{n-1,\star}) \rangle_{hu}\} \\ &= \min \{k, \langle \mathbf{o}h_{j-n,\star} \rangle_{hu} + \langle \mathbf{o}l_{n-1,\star} \rangle_{hu}\} \\ &= \min \{k, \text{ones}(\mathbf{a}_{j..n}) + \text{ones}(\mathbf{a}_{n-1..0})\} \\ &= \min \{k, \text{ones}(\mathbf{a}_{j..0})\} \\ &= O_j \end{aligned}$$

C.2.4 Different Interpretation

The following two lemmas satisfy a column-wise interpretation of the output array $\mathbf{o}_{\star,\star}$. This way, a FFk1hu_n can be used to actually *find* the first k ones, i.e. mark up their positions.

Lemma C.1 *Let $\mathbf{o}_{\star,\star} \in \{0,1\}^{k \cdot n}$ satisfy $\langle \mathbf{o}_{j,\star} \rangle_{hu} = O_j$. Then:*

$$O_{j-1} = O_j - 1 = l \implies \langle \overline{\mathbf{o}_{\star,l}} \rangle_{hu} = j$$

Proof. Assume $\text{ones}(a_{(j-1) \cdot 0}) = \text{ones}(a_{j \cdot 0}) - 1 = l$. Because of half-unary number encoding, we have $o_{j-1,1} = 0 \wedge o_{j,1} = 1$. Since the $\text{ones}(\cdot)$ function is monotonous in growing inputs, i.e. $\text{ones}(x_{m \cdot 0}) \geq \text{ones}(x_{m-1 \cdot 0})$, the l -th column of $\mathbf{o}_{*,*}$ satisfies:

$$o_{j',l} = \begin{cases} 0 & \text{for } j' \in \{0, \dots, j-1\} \\ 1 & \text{for } j' \in \{j, \dots, n-1\} \end{cases}$$

Therefore $\langle \overline{o_{*,l}} \rangle_{hu} = j$.

Lemma C.2 *Let $\mathbf{o}_{*,*} \in \{0, 1\}^{k \cdot n}$ satisfy $\langle o_{j,*} \rangle_{hu} = O_j$. Then:*

$$\text{ones}(a_{n-1 \cdot 0}) = l \implies \forall l' \geq l : \langle \overline{o_{*,l'}} \rangle_{hu} = n$$

Proof. With $\langle o_{n-1,*} \rangle_{hu} = \text{ones}(a_{n-1 \cdot 0}) = l$ we obtain $\mathbf{o}_{n-1,k-1 \cdot 1} = 0^{k-l}$. Since $\text{ones}(x_{m \cdot 0}) \geq \text{ones}(x_{m-1 \cdot 0})$, this results in the claim $\mathbf{o}_{*,k-1 \cdot 1} = 0^{(k-l)n}$.

Corollary C.3 (Sorting) *Let $K := O_{n-1}$. Define*

$$s_l := \langle o_{*,l} \rangle_{hu} \quad \text{for } l \in \{0, \dots, K-1\}$$

Then (s_1, \dots, s_{K-1}) is the prefix of length K of the sorted sequences of indices of the inputs bits being 1. Formally:

$$s_l < s_{l+1} \quad \wedge \quad i_{s_{l+1} \dots s_1} = (1, 0^{s_{l+1} - s_l - 1}, 1)$$

Proof. The corollary follows from lemmas C.1 and C.2.

C.3 Multiple Incrementer

The circuit $K\text{-MInc}_n$, $K = 2^k$ and $k, n \in \mathbb{N}$, computes the following function:

$$\begin{aligned} K\text{-minc}_n : \{0, 1\}^n &\longrightarrow \{0, 1\}^{(K+1) \cdot n} \\ \mathbf{a}_* &\longmapsto \mathbf{o}_{*,*} \end{aligned}$$

with

$$\langle \mathbf{o}_{K',*} \rangle_2 \equiv_{2^n} \langle \mathbf{a}_* \rangle_2 + K'$$

We develop an implementation suitable for small k . The computation of the outputs proceeds in two steps. The input string \mathbf{a}_* is split into a high and a low part:

$$\begin{aligned} \mathbf{al}_* &:= \mathbf{a}_{k-1 \cdot 0} \\ \mathbf{ah}_* &:= \mathbf{a}_{n-1 \cdot k} \end{aligned}$$

To generate the low parts of the results $\mathbf{ol}_{*,*}$ we use a $2K(k+1)$ -bit barrel right-shifter SH. We let the shifter operate on $2K$ groups of $(k+1)$ bits length. It receives \mathbf{al}_* as shift distance and its inputs are the bit string of length $(k+1)$ in lexicographical order:

$$\begin{aligned} \text{SH.input}_{i,*} &:= \text{bin}(i) \\ \text{SH.dist}_* &:= \mathbf{al}_* \\ \mathbf{ol}_{i,*} &:= \text{SH.output}_{i,*} \end{aligned}$$

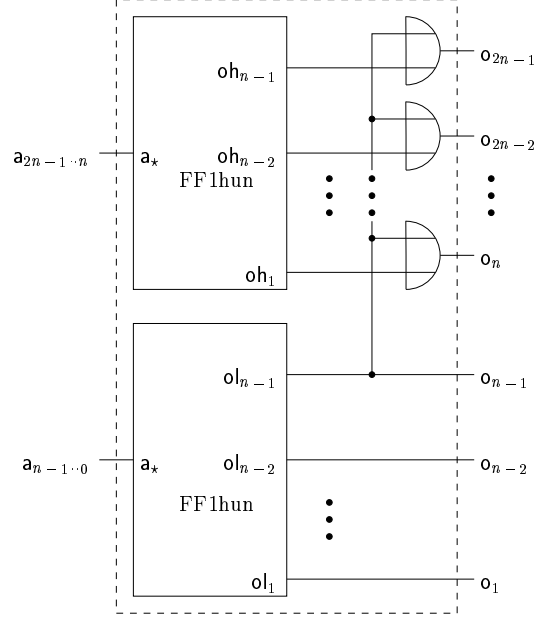


Figure C.1: Definition of a find-first-one half-unary circuit

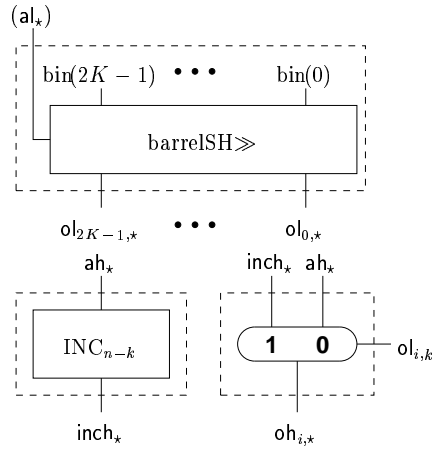


Figure C.2: The multiple incrementer

Thereby the shifter computes

$$\begin{aligned} \langle \text{ol}_{i,\star} \rangle_2 &= \langle \text{bin}(i) + \langle \text{al}_\star \rangle_2 \mod 2K \rangle_2 \\ &\equiv_{2K} i + \langle \text{al}_\star \rangle_2 \end{aligned}$$

The uppermost bit $\text{ol}_{i,k}$ can be used to select between the original and an incremented version inch_\star , $\langle \text{inch}_\star \rangle_2 \equiv_{2^{n-k}} \langle \text{ah}_\star \rangle_2 + 1$, of the input bus in conditional-sum-adder fashion:

$$\text{oh}_{i,\star} := (\text{ol}_{i,k} ? \text{ah}_\star : \text{inch}_\star)$$

The output busses are composed as

$$\text{o}_{i,\star} := (\text{oh}_{i,\star}, \text{ol}_{i,k-1} \dots 0)$$

Bibliography

- [AKT86] R. D. Acosta, J. Kjelstrup, and H. C. Torng. An instruction issuing approach to enhancing performance in multiple function unit processors. In *IEEE Transactions on Computers*, volume C-35, pages 815–828. 1986.
- [AMS97] T.N. Vijaykumar A. Moshovos, S. E. Breach and G. S. Sohi. Dynamic speculation and synchronization of data dependences. In *Proceedings of the 24th Annual International Symposium on Computer Architecture Conference*, 1997.
- [AS95] Don Anderson and Tom Shanley. *Pentium Processor System Architecture, second edition*. Addison-Wesley, Amsterdam;Sydney;Singapore, 1995.
- [Bha96] Dileep P. Bhandarkar. *Alpha Implementations and Architecture, Complete Reference and Guide*. Digital Press, Boston;Oxford;Tokyo, 1996.
- [Del98] Peter Dell. Die Auswirkung von Mechanismen zur Out-of-Order Ausführung auf den Cyclecount von RISC-Architekturen. Master’s thesis, Universität des Saarlandes, FB. Informatik, 7/ 1998.
- [Ger98] Nikolaus Gerteis. Die Auswirkung von Mechanismen für die präzise Interruptbehandlung auf den Cyclecount von RISC-Prozessoren. Master’s thesis, Universität des Saarlandes, FB. Informatik, 4/ 1998.
- [Grü94] Thomas Grün. *Quantitative Analyse von I/O-Architekturen*. PhD thesis, Universität des Saarlandes, FB. Informatik, 1994.
- [Joh91] Mike Johnson. *Superscalar Microprocessor Design*. Prentice-Hall, Englewood Cliffs, New Jersey, 1991.
- [KP95] Jörg Keller and Wolfgang J. Paul. *Hardware design – Formaler Entwurf digitaler Schaltungen*. Stuttgart, Leipzig, 1995.
- [Krö99] Daniel Kröning. Design and evaluation of a RISC processor with a Tomasulo scheduler. Master’s thesis, University of Saarland, Computer Science Department, Germany, 1999.
- [Lei99] Holger Leister. *Quantitative Analysis of Precise Interrupt Mechanisms for Processors with Out-Of-Order Execution*. PhD thesis, University of Saarland, Computer Science Department, Germany, 1999.
- [LS84] Johnny K. F. Lee and Alan J. Smith. Branch prediction strategies and branch target buffer design. *Computer*, 17(1):6–22, January 1984.
- [MP95] Silvia M. Müller and Wolfgang J. Paul. *The complexity of simple computer architectures*, volume 995 of *Lecture Notes in Computer Science*. Springer-Verlag Inc., New York, NY, USA, 1995.

- [MP00] Silvia M. Müller and Wolfgang J. Paul. Computer Architectures: Complexity and Correctness. Draft. EMail: {smueller,wjp}@cs.uni-sb.de, 2000.
- [Mül97] Silvia M. Müller. Vorlesung Rechnerarchitektur II, WS 96/96, Universität des Saarlandes, FB. Informatik. 1997.
- [PS98] Wolfgang J. Paul and Peter-Michael Seidel. On the complexity of Booth recoding. In *Proc. 3rd Conference on Real Numbers and Computers (RNC3)*, pages 199–218, 1998.
- [SP88] James E. Smith and Andrew R. Pleszkun. Implementing precise interrupts in pipelined processors. In *IEEE Transactions on Computers*, volume C-37, pages 562–573. 1988.
- [Tom67] Robert M. Tomasulo. An efficient algorithm for exploiting multiple arithmetic units. In *IBM Journal of Research and Development*, volume 11 (1), pages 25–33. IBM, 1967.
- [Weg87] Ingo Wegener. *The Complexity of Boolean Functions*. Stuttgart, 1987.
- [WS84] Shlomo Weiss and James E. Smith. Instruction issue logic for pipelined supercomputers. In *Proceedings of the 11th Annual International Symposium on Computer Architecture*, pages 110–118, Ann Arbor, Michigan, June 5–7, 1984. IEEE Computer Society and ACM SIGARCH.
- [WS94] Shlomo Weiss and James E. Smith. *Power and PowerPC*. Morgan Kaufmann, San Francisco, 1994.
- [Yeh93] Tse-Yu Yeh. *Two-Level Adaptive Branch Prediction and Instruction Fetch Mechanism for High Performance Superscalar Processors*. PhD thesis, University of Michigan, USA, 1993.