An Enhancement for a Scheduling Logic Pipelined over two Cycles

Rubén Gran, Enric Morancho, Àngel Olivé & José María Llabería
Department of Computer Architecture, Polytechnic University of Catalonia.
{rgran, enricm, angel & llaberia}@ac.upc.edu

Abstract—**Out of order processors use the dynamic scheduling logic to expose and exploit parallelism. Pipelining this logic may sacrifice the ability to execute dependent instructions in consecutive cycles. Several previous studies have shown that pipelining the scheduling logic over two cycles degrades performance; our evaluations, in a 4-way machine, on SPEC-2000 integer benchmarks show a performance degradation about 11% with respect to an unpipelined scheduling logic.**

In this work, we present two non-speculative enhancements for a scheduling logic pipelined over two cycles. The idea is computing in advance which instructions will be woken-up by all instructions that are currently competing for selection. Once all of them have been selected, the pre-computed group of instructions can compete for selection in next cycle. The enhancement goal is tolerating the scheduling-loop latency when not enough ILP is available, through the scheduling of dependent instructions in consecutive cycles.

Our results in a 4-way machine show that our two proposed enhancements perform, in average, slightly better than two previously proposed speculative schedulers. The performance of our proposals is within a 2.6% and 2% of an unpipelined ideal scheduler, respectively.

Keywords—**Dynamic scheduling, Pipelined scheduling logic, Back-to-back execution**

I. Introduction

The dynamic scheduling logic allows both exposing and exploiting the instruction-level parallelism (ILP). The scheduling task is divided into two phases: wakeup and select. The wakeup logic marks instructions as ready when their data dependencies are satisfied. The select logic picks instructions for execution from the pool of ready instructions by considering instruction priorities and available resources.

Both the wakeup logic and the select logic form a hardware loop, the scheduling loop, because an instruction cannot be scheduled until its producer instructions have been scheduled. Assuming that the producer-instruction latency is one cycle then, in order to execute dependent instructions in consecutive cycles, the scheduling task must be performed in one cycle. A producer instruction and its consumer instruction are executed back-to-back when the consumer instruction consumes the produced result as soon as it is available.

Enlarging the issue queue to expose more ILP may increase the latency needed to wakeup and select instructions, which may require reducing clock frequency. An approach to either maintaining or increasing clock frequency is pipelining the scheduling logic over several cycles, but then the IPC may decrease, because the scheduling logic sometimes is unable to issue dependent instructions in consecutive cycles. Our experimental results with SPEC-2000 integer benchmarks in a 4-issue machine show that pipelining the scheduling logic over two cycles has an average IPC degradation about 11% with respect to an unpipelined scheduling logic. Other authors report similar results ([2], [16], [20]).

Techniques that allow pipelining the scheduling logic without sacrificing the back-to-back execution of dependent instructions are an option to design high-frequency processors. However, proposed techniques ([2], [20]) are speculative.

In this paper, we enhance a scheduling logic pipelined over two cycles to increase its performance. The proposed enhancement is non-speculative and able to execute dependent instructions in consecutive cycles when not enough ILP is available. Consequently, we manage to tolerate the scheduling-logic latency. The idea of the enhancement is computing in advance which instruction group will be woken up by all one-cycle execution-latency instructions that are currently competing for selection. Then, once all these instructions have been selected, the pre-computed instruction group can compete for selection in next cycle and back-to-back execution may be performed.

Our results show that our two proposed enhancements perform, in average, better than two previously proposed speculative schedulers ([2], [20]), in SPEC-2000 integer benchmarks. The performance of our proposal is within a *2.6%* and *2%* of an ideal scheduler (unpipelined), respectively.

This paper is structured as follows: Section II outlines the processor model being used and motivates the work. Section III describes the proposed enhancement. Section IV details the simulation environment. Section V evaluates the proposed models and compares it with two previously proposed mechanisms. Section VI discusses related work and Section VII concludes this paper.

II. Baseline processor model

The pipeline of a dynamically scheduled processor is shown in Figure 1, where each stage can take one or more cycles.

In the front-end stages of the pipeline (fetch, decode and rename stages), instructions are brought from the instruction cache, decoded and false register dependencies are removed. After that, the instructions are dispatched into the issue queue and wait there for the availability of both their source operands and execution resources. When an instruction is selected for execution, the payload and its source registers are read in following cycles. With its source operands, the instruction is executed and its result is written into the register file. Finally, the instruction waits for committing in program order.

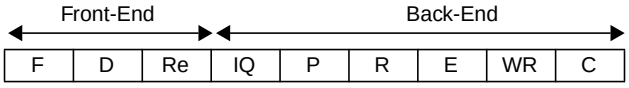

Fig. 1. Processor Pipeline. F: Fetch, D: Decode, Re: Rename, IQ: Issue Queue, P: Read Payload, R: Read Register File, E: Execution; WR: Write Register File, C: Commit.

Wakeup logic. We use a wired-OR style wakeup logic array ([2], [8]). Dependencies are indicated using an *instructions-instructions*[1] wakeup matrix [2] (or *physical registers-instructions*[2] wakeup matrix [8]). Bit vectors (rows) perform dependence tracking. Each bit in the vector represents the dependence on a parent instruction [2] (or on the data availability of a physical register [8]). When an instruction is issued, it sets the wakeup line (column) corresponding to its own issue-queue entry [2] (or to its destination physical register [8]). Each instruction monitors the readiness of its source operands every cycle by checking if all wakeup lines of matching dependence bits are set. Each issue-queue entry corresponding to a ready instruction activates a request signal in order to notify its readiness.

Select logic. The input of the select logic are request signals from the wakeup logic plus priority information. The select logic picks the oldest ready instructions taking into account available resources at each issue port. Instructions selected by the select logic are the input of the wakeup logic on next clock cycle in order to wakeup instructions dependent on the selected ones.

Figure 2 shows diagrams of both one-cycle latency (unpipelined) and two-cycle latency scheduling loops. As a general rule, back-to-back execution is possible only if the execution latency of the producer instruction is greater than or equal to the scheduling-loop latency. In [2], [16], [20], their authors have concluded that back-to-back execution is a performance goal.

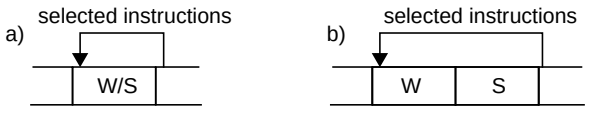

Fig. 2. Diagrams of scheduling loops. a) one-cycle latency, b) two-cycle latency. (W: Wakeup, S: Select)

Table I shows the distribution of committed instructions in the SPEC-2000 benchmarks taking into account their execution latency and if the instructions produce a value that is stored in the register file (Section IV details benchmarks, simulated intervals and the execution latency of the instructions). We observe that integer benchmarks double the amount of one-cycle execution-latency instructions with respect to floating-point benchmarks; consequently, integer benchmarks will be more sensitive to the scheduling-loop latency.

TABLE I DISTRIBUTION OF COMMITTED INSTRUCTIONS ACCORDING TO THEIR EXECUTION LATENCY IN THE SPEC-2000 BENCHMARKS.

Benchmarks	updating register file: execution latency: one cycle	updating register file: execution latency: multicycle	not updating register file
Integer	44.30%	32.05%	23.65%
Floating Point	23.62%	62.40%	13.98%

In this paper, the baseline processor has a two-cycle latency scheduling loop. Then, at least, there is a two-cycle delay between issuing an instruction and issuing its dependent instructions. So, in the issue cycle between issuing an one-cycle execution-latency instruction and issuing its dependent instruction, the scheduling logic must be able to exploit ILP in order not to degrade performance with respect to the unpipelined scheduling logic. For multi-cycle execution-latency producer instructions (greater than one cycle), pipelining the scheduling logic does not degrade performance with respect to an unpipelined scheduling logic.

III. ENHANCED SCHEDULING LOGIC

In this section we describe a non-speculative enhancement that improves the performance of a two-cycle latency scheduling logic.

A. Base enhancement (E)

The idea is computing in advance which instructions will be woken up by all instructions currently competing for selection. Once, all of them have been selected, the pre-computed group of instructions can compete for selection in next cycle. This idea is applied only to instructions dependent on instructions with an execution latency shorter than the scheduling-loop latency. Therefore, despite of the scheduling-logic latency, back-to-back execution is possible. The remaining instructions use the conventional two-cycle latency scheduling logic.

Figure 3 shows a scheme of a two cycle latency scheduling logic with the proposed enhancement logic. The base two-cycle latency scheduling logic is composed by the wakeup matrix A and the select logic. The output of the selection logic is, in the next cycle, the input of the wakeup matrix A. On each selection cycle, the select logic picks the oldest instructions that remain in its input, taking into account resource availability at

1 Each wakeup-matrix row and each wakeup-matrix column corresponds to an instruction inserted into the issue queue.

2 Each wakeup-matrix row corresponds to an instruction inserted into the issue queue and each wakeup-matrix column corresponds to a physical register.

each issue port.

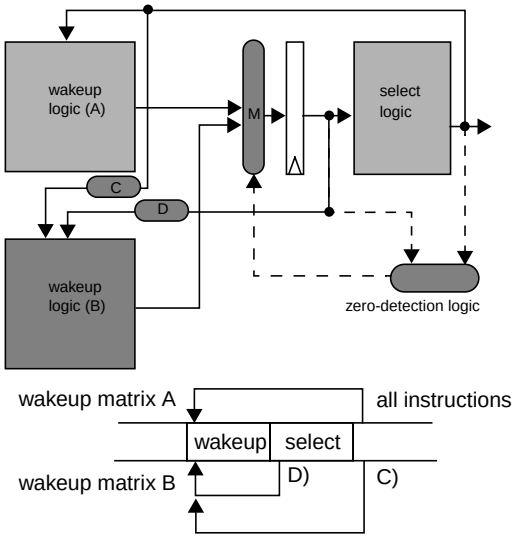

Fig. 3. Base proposal. The output of logic C are the instructions with execution latency greater than or equal to scheduling-logic latency that have been selected by the selection logic. The output of logic D are the instructions with execution latency shorter than the scheduling-logic latency that are currently competing for selection. The logic labeled as M merges request signals of the wakeup matrix B with request signals of the wakeup matrix A. Also, the M logic eliminates duplicated request signals. The zero-detection logic detects if all one-cycle execution-latency instructions, that are currently competing for selection, have been scheduled.

The wakeup matrix B computes in advance which dependent instructions will be woken up by the one-cycle execution-latency instructions which are in the input of the selection logic.

In rename stage, each instruction is classified taking into account the latency of the instructions that wake it up (parent instructions). At dispatch time, all instructions are stored in wakeup matrix A. In wakeup matrix B are stored instructions that may be woken up by instructions whose execution latency is shorter than the scheduling-logic latency.

Inputs of the wakeup matrix B are: a) one-cycle execution-latency instructions in the input of the select logic and b) instructions selected by the selection logic that have an execution latency greater than or equal to the scheduling-loop latency. These inputs are respectively calculated by filters D and C (Figure 3), using the instruction classification performed in rename stage.

The logic circuit M merges the request signals that are generated by both wakeup matrices. These request signals are merged if all one-cycle execution-latency instructions pending for selection in the input of the selection logic have been scheduled (zero-detection logic).

The request signal of an instruction that is woken up by an one-cycle execution-latency instruction is activated in both wakeup matrices at different cycles. First, when the parent instruction is in the input of the select logic, the request signal is activated in wakeup matrix B. Later, when the parent is selected, the request signal is activated in wakeup matrix A. Then, for every instruction inserted in both wakeup matrices, the logic circuit M filters out the latest request signal to arrive related to the same entry. Therefore, only one request signal is observed by the selection logic. Another case happens when the latest arriving operand of an instruction stored in wakeup matrix B is produced by an instruction whose execution latency is greater than or equal to the scheduler-logic latency. In this case both request signals are activated in the same cycle.

Figure 4 shows an example of instruction scheduling in which, on behalf of understanding, only one instruction can be issued per cycle. The IQ label means that the instruction is waiting to be ready in the issue queue. A and B labels mean that the instruction wakes up in wakeup matrices A and B respectively. The RI label symbolizes that the instruction is waiting for selection in the input of the selection logic. Finally, the S label means that the instruction is selected for execution.

Cycles	1	2	3	4	5
1. add r1←r2, r3	A	S			
2. add r4←r5, r6	A	RI	S		
3. sub r9←r1, r7	IQ	B	A	S	
4. sub r10←r9, r8	IQ	IQ	IQ	B	A/S

Fig. 4. Scheduling example of the proposed mechanism. A bar between cycles indicates that the output of wakeup matrix B is merged with the output of wakeup matrix A.

In Figure 4, the source operands of the first two instructions are available on dispatch, and consequently, both instructions wake up in cycle 1. Next cycle, instructions 1 and 2 are candidate for selection. Also in cycle 2, instruction 1 wakes instruction 3 up in wakeup matrix B and it waits to be merged until all one-cycle execution-latency instructions in the input of the select logic have been scheduled (instructions 1 and 2). At the end of cycle 3, the output of the wakeup matrix B is merged with the output of the wakeup matrix A, and also, logic circuit M detects the duplication of request signals of the instruction 3. In cycle 4, instruction 3 wakes instruction 4 up in wakeup matrix B, and also, instruction 3 is selected by the selection logic. At the end of cycle 4, the zero-detection logic detects that there is no instruction pending to be selected and the output of the wakeup matrix B is merged again. Then, the instruction 4 is in the input of the select logic and it is selected in cycle 5. In cycle 5, the logic circuit M detects duplication of the request signal of instruction 4.

B. Adding instruction fusing (E-F)

The proposed enhancement can be improved by taking advantage of a program characteristic: a large number of instructions has only one source operand (avg: *78.6%* of committed instructions in SPECInt 2000). Moreover, at dispatch time, some two-operand instructions have already available one of them. Consequently only one operand must be tracked by the wakeup matrix.

Then, we make use of fusing instructions (a producer instruction and its dependent one) in order to favour back-to-back execution of dependent instructions. Two instructions are fused when producer instruction is an one-cycle execution-latency instruction and the consumer instruction only depends on this instruction.

In our model E-F, the advantages of fusing instructions

are twofold. First, the consumer instruction can compete for selection once the producer instruction have been selected. Therefore, back-to-back execution is possible. Second, it is not necessary to store the consumer instruction in the wakeup matrix B for waking it up.

The possibility of fusion is detected in rename stage. The fused instructions must belong to the same dynamic basic block and the producer instruction must be in the issue queue.

The instruction fused with its producer instruction is the first one in program order that satisfies the previous conditions. Note that the proposed instruction fusing is simple because both producer and consumer instructions belong to the same basic block.

In our evaluations, two issue-queue entries are allocated to fused instructions in wakeup matrix A and two issue cycles are needed to schedule them. Therefore, we maintain the same pressure than previous models over the issue-queue entries of wakeup matrix A and the issue ports.

IV. Simulation Environment

A. Processor model

We have modified SimpleScalar 3.0c [2] in order to model a Reorder Buffer and separate issue queues (IQ). We assume an out-of-order processor with fifteen stages from Fetch to IQ and two stages between IQ and Execution. Other processor and memory parameters are listed in Table II.

TABLE II Processor and memory parameters

	Model
Fetch and Decode width	4
Branch predictor: hybrid (bimodal, gshare)	16 bits
ROB size	128 entries
LSQ size	64 entries
Issue-queue size Integer / Floating point	32 / 20 entries
Functional Units Integer / Floating point	4 / 2
Memory access ports	2

Memory hierarchy	
L1 I-cache	32KB, 4-way, 2 cycles
L1 D-cache	32KB, 4-way, 2 cycles
Line size	32 B
L2 Unified Cache	256 KB, 4-way, 12 cycles
Line size	32 B
L2-Main memory bus	8B / 2 cycles
Main memory latency	100 cycles

In Table III are listed the instructions latencies assumed in this work.

TABLE III Execution latency (in cycles) of the instructions.

	Latency		Latency
ALU	1	Floating point add, mul	4 pipelined
Load	3	Floating point divide	15 not pipelined
integer multiply	10 not pipelined	others	1

We split store instructions into two instructions: STA (store address computation) and STD (store data). Therefore, two issue-queue entries are allocated to each store instruction.

A load instruction can be issued only after issuing all the STA instructions corresponding to the store instructions older than the load instruction. Consequently, we made each load instruction dependent on all its older STA instructions.

The IQ is divided into an integer and floating-point IQ's. Our proposals are applied only to the integer IQ because the execution latency of most FP instructions is greater than the scheduling-loop latency.

B. Workload

We use SPEC2000 integer benchmarks compiled with full optimizations on a Alpha machine. We simulate a contiguous run of 100M-instruction from SimPoints [19] after a warming-up of 100M-instruction. Table IV shows their input data sets.

TABLE IV Simulated benchmarks and their input data set.

Bench.	Data set	Bench.	Data set	Bench.	Data set
bzip2	program-ref	gzip	program-ref	twolf	ref
crafty	ref	mcf	ref	vortex	one-ref
eon	rushmeier-ref	parser	ref	vpr	route-ref
gcc	166-ref	perl	diffmail-ref		

V. Results

To evaluate the performance of our proposed enhancement we have simulated several models with a two-cycle latency scheduling loop.

- A baseline model (B) where back-to-back execution of dependent instructions is sacrificed when producer instructions have one-cycle execution latency. Also, we model instruction fusing (B-F) with same conditions than in Section B.
- Two models implementing our proposals: E and E-F.
- For comparison purposes, B-Double model doubles the number of integer issue-queue entries of the baseline model. This model is intended for showing us what it is more cost-effective, dedicating added IQ entries to either expose more parallelism or favour the back-to-back execution of dependent instructions.
- For comparison purposes, we model the Speculative Wakeup (SW, [20]) and the Select-Free (SF, [2]) mechanisms. Both of them are speculative mechanisms designed to tolerate the scheduling-logic latency. They are described in Section VI. In our evaluations, the Speculative Wakeup mechanism is implemented by using two wakeup matrices (the first one for tracking the parent instructions and the second one for tracking the grandparent instructions). In the Select-Free mechanism, speculation is checked in register-read stage.

Moreover, we simulate an ideal model (ID) with unpipelined scheduling-loop (that is, its latency is one cycle). However, in order to remove the effect of a branch-misprediction penalty shorter than in the other models, its pipeline depth is kept consistent with them by adding one extra stage in the front-end.

Table V shows the IPC achieved for each benchmark

in B model.

TABLE V IPC OF THE BASELINE MODEL (B).

bzip2	crafty	eon	gap	gcc	gzip	mcf	parser	perl	twolf	vortex	vpr
1.34	1.86	2.12	1.94	1.38	1.60	0.13	1.01	1.38	0.88	2.38	0.77

Figure 5 shows the speed-up of all models with respect to the B model. We present individual results for each SPEC-2000 integer benchmark and two average values: for all benchmarks (HM) and for all benchmarks but *mcf* (HM-mcf) due to its biased memory behaviour.

Our proposed models, in average, outperform both the B-Double and the B-F models. The B-Double model performs better than our proposed models in benchmarks *mcf* and *vortex*. However, in the other benchmarks, doubling the number of issue-queue entries to expose more parallelism is not cost-effective. It is better to favour the back-to-back execution of dependent instructions. Our proposed models also perform better than B-F model in all benchmarks. B-F model increases performance, in average, about *2.8%* with respect to B model. However, in *gzip*, performance improvement reaches *9.1%*.

We observe that our proposed models, in average, outperform the speculative models (SW and SF). The SW model performs better than our proposed models only in benchmark *vpr*; the SF model performs better than E model in benchmarks *bzip2, gzip, perl and vpr*. However, the E-F model is outperformed by the SF model only in the benchmark *vpr.*

Performance of E and E-F models are, in average, within *2.6%* and *2%* of the ID model, respectively. In the E-F model, instruction fusing permits a consumer instruction, that has been fused with its producer, to avoid waiting for the next merge operation to compete for selection. Otherwise, in the E model, consumer instructions have to wait for the complete schedule of all one-cycle execution-latency instructions, that are currently competing for selection. Therefore, in the E-F model, those instructions could save some cycles to reach the selection stage. In average, a *19.2%* of dispatched instructions are fused with their producer instruction.

While both the SF and the SW models are speculative, our proposed models are not. The SF model must re-schedule some instructions, that have been speculatively woken-up. This involves activity, which wastes energy, in both the select logic and the register file. Our evaluations show that, in average, in the SF model the re-schedulings affect to a 3.4% of committed instructions and a 3.0% of selections by the select logic.

The SW model may select instructions whose selection will be later nullified because parent instructions have not been issued. These *false selections* affect, in average, to a 7.6% of the committed instructions and a 4.6% of selections by the select logic.

The SW model and our proposed models use two wakeup matrices. In the SW model, all instructions are stored in both wakeup matrices. However, in our proposed models, the wakeup matrix B stores bit dependence vectors of an instruction only if it can be woken up by an one-cycle execution-latency instruction. In the proposed E model the average occupancy of wakeup matrix B is a 19% smaller than the average occupancy of wakeup matrix A. And for the model E-F, wakeup matrix B is occupied, in average, a 34% less than the wakeup matrix A.

In the SW model and in our proposed models, the empty entries of both wakeup matrices can be dynamically deactivated [1]. Then, our proposed models are more energy-efficient than SW model because the average occupancy of wakeup matrix B is smaller. Moreover, once the request signal of an instruction is detected by logic circuit M, its mirror entry in the other wakeup matrix can be deactivated.

VI. RELATED WORK

In order to reduce the scheduling latency, Palacharla et al. [16] proposed dispatching chains of dependent instructions into FIFO queues; the instructions considered to be issued are only the instructions heading each FIFO queue. Another works preschedule the instructions taking advantage of the fact that most instruction latencies are known at decode time ([5], [7], [14]). At dispatch time, instructions are sorted into a buffer according to their predicted issue cycle. The schemes mainly differ in the mechanism that deals with variable-latency instructions, e.g. load instructions, and their chains of dependent instructions; a structure like an issue queue is used for these cases. All these techniques require estimating the issue cycle of instructions before inserting them in the buffer structure.

Some proposals exploit the fact that most register-writing instructions have, at most, one dependent instruction currently in the issue queue. Based

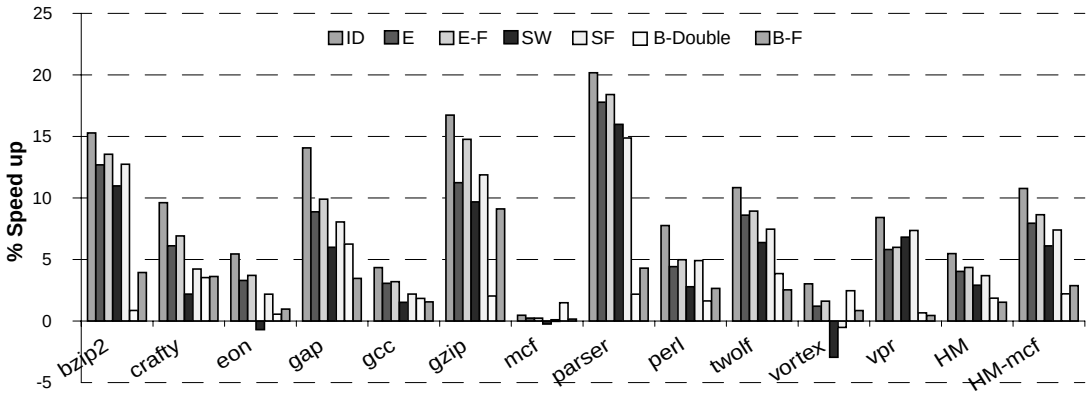

Fig. 5. Speed-up with respect to the B model

on this observation, the proposed designs have structures that keep track of one or several instructions that consume a produced register value ([5], [21]). These techniques require additional hardware support for branch-misprediction recovery unless the recovery is initiated only when the branch instruction becomes the oldest instruction in flight. Other proposal uses RAM bitmap arrays to identify all the successors of each instruction in the issue queue [9]. A new design that reduces the area cost for large issue queues was proposed by K. Hsiao and C. Chen in [10].

The observation that many instructions already have one or two ready source operands at dispatch time has been used to reduce the load capacitance of the wakeup tag bus in schedulers that use CAM schemes to wakeup; consequently, the wakeup latency may be reduced ([6], [9], [13]).

I.Kim and M. Lipasti proposed a hardware mechanism that dynamically detects dependent pairs of instructions and fuses them in order to be scheduled together [12]. So, the scheduling-loop latency (two cycles) is hidden because the scheduling granularity has been increased. A later work removes complexity from hardware and enables more sophisticated fusing heuristics using dynamic-translation software that becomes part of the processor design [11]. Other related works use intensive hardware to combine dependent operations [18] that are issued speculatively or need static compiler support [3].

Several works use speculation to break the scheduling loop. Stark et al. [20] proposed speculatively waking instructions up by their grandparents. This proposal allows pipelining the scheduling loop over two cycles. The speculative wakeup of an instruction is confirmed after their parents are selected. A false selected instruction affects performance only if it prevents really ready instructions from being selected for execution. Brown et al. [2] proposed a speculative technique, named Select-Free, which moves the selection logic off the critical loop; this allows the scheduling loop to take just one cycle. The technique allows all woken-up instructions broadcasting the tags into the issue queue in the following cycle, even though some of them may have not been selected for execution yet. Contention for issue ports can produce the misspeculated wakeup of a chain of dependent instructions. Therefore, the availability of source operands of each issued instruction are checked before execution stage. Both proposals allow back-to-back scheduling of dependent instructions. Focusing in the SW model, we have measured, in average, a *7.6%* of committed instructions that are falsely selected, which unnecessarily utilize the selection logic. In the Select-Free model happens something similar. In this case, measures show, in average, a *3.4%* of committed instructions that have to be re-scheduled. Obviously, re-scheduled instructions incurs in a unnecessary utilization of both the selection logic and the register file.

VII. Conclusions

In this paper, we have proposed two non-speculative enhancements (E and E-F) for a scheduling logic pipelined over two cycles. These enhancements try to tolerate the latency of the scheduling loop when there is not enough available ILP. In order to do this, our proposals compute in advance which instructions will be woken up by all one-cycle execution-latency instructions that are currently competing for selection. This pre-computed group of instructions can compete for selection, once that all previous one-cycle execution-latency instructions, which are currently competing for selection, have been selected. Moreover, we have improved the base enhancement (E) using instruction fusing (E-F).

Our evaluations have shown that E and E-F models perform, in average, within a *2%* and *2.6%* of an ideal (unpipelined) scheduler, respectively. In comparison with baseline model B (scheduling loop pipelined over two cycles), E and E-F increase, in average, performance a *7,9%* and a *8.6%*, respectively. Also, E and E-F perform, in average, slightly better than the two previously-proposed speculative schedulers (Select Free and Speculative Wakeup).

References

[1] Y. Bai; R. I. Bahar. A dynamically reconfigurable mixed in-order/out-of-order issue queue for power-aware microprocessors, Annual Symposium on VLSI, p. 139-146, 2003.

[2] M. Brown et al. Select-Free Instruction Scheduling Logic. Micro-2001, p. 204-213.

[3] A. Bracy et al. Dataflow Mini-Graphs: Amplifying Superscalar Capacity and Bandwidth. Micro-2004, p.18-29.

[4] D.C. Burger and T.M. Austin, “The SimpleScalar Tool Set, Version 2.0,” UW Madison Computer Science T. R. #1342, June 1997.

[5] R. Canal and A. González. A Low-Complexity Issue Logic. ICS-2000, p. 327-335.

[6] D. Ernst and T.M. Austin. Efficient dynamic scheduling through tag elimination. ISCA-2002, p.37-46

[7] D. Ernst et al. Cyclone: A Broadcast-Free Dynamic Instruction Scheduler with Selective Replay. ISCA-2003, p. 253-262.

[8] J.A. Farrel and T.C Fischer. Issue Logic for a 600 Mhz Out-of-order Execution Microprocessor. Journal of Solid-State Circuits, Vol 33(5), pp 707-712, 1998.

[9] M. Goshima et al.. A high-speed dynamic instruction scheduling scheme for superscalar processors. Micro-2001, p. 225-236.

[10] K. S. Hsiao and C.H. Chen. An Efficient Wakeup Design for Energy Reduction in High-Performance Superscalar Processors. Conf. on Computing Frontiers, 2005, p. 353-360.

[11] S. Hu et al. An Approach for Implementing Efficient Superscalar CISC Processors. HPCA-2006, p. 40-51.

[12] I. Kim and M. H. Lipasti, Macro-op Scheduling: Relaxing Scheduling Loop Constraints. Micro-2003, p. 277-288.

[13] I. Kim and M. H. Lipasti. Half-Price Architecture. Int. Symp. on Computer Architecture. 2003, p. 28-38.

[14] P. Michaud and A. Seznec. Data-flow prescheduling for large instruction windows in out-of-order processors. HPCA-2001, p. 27-36.

[15] S. Önder and R. Gupta. Instruction Wake-up in Wide issue superscalars. European Conf. on Parallel Processing, 2001, p. 418-427.

[16] S.Palacharla et al. Quantifying the complexity of superscalar processors. T.R. University of Wisconsin-Madison. Nov 1996.

[17] E. Perelman et al.. Picking Statistically Valid and Early Simulation Points. PACT-2003, p. 244-255.

[18] P. G. Sassone and D. Scott Wills. Dynamic Strands: Collapsing Speculative Dependence Chains for Reducing Pipeline Communication. Micro-2004, p. 7-17

[19] T. Sherwood et al., “Automatically Characterizing Large Scale Program Behaviour,” in Proc. of ASPLOS, Oct. 2002, p. 45-57.

[20] J. Stark et al. On pipelining dynamic instruction scheduling logic. Micro-2000, p. 204-213.

[21] S. Weiss and J.E. Smith. Instruction issue logic in pipelined supercomputers. IEEE Trans. on Computers, 33: p.1013-1022, Nov.84.