

FREESS: An Educational Simulator of a RISC-V-Inspired Superscalar Processor Based on Tomasulo's Algorithm

Roberto Giorgi
giorgi@unisi.it
University of Siena
Siena, Italy

Abstract

FREES is a free, interactive simulator that illustrates instruction-level parallelism in a RISC-V-inspired superscalar processor. Based on an extended version of Tomasulo’s algorithm, FREES is intended as a hands-on educational tool for Advanced Computer Architecture courses. It enables students to explore dynamic, out-of-order instruction execution, emphasizing how instructions are issued as soon as their operands become available.

The simulator models key microarchitectural components, including the Instruction Window (IW), Reorder Buffer (ROB), Register Map (RM), Free Pool (FP), and Load/Store Queues. FREES allows users to dynamically configure runtime parameters, such as the superscalar issue width, functional unit types and latencies, and the sizes of architectural buffers and queues.

To simplify learning, the simulator uses a minimal instruction set inspired by RISC-V (ADD, ADDI, BEQ, BNE, LW, MUL, SW) which is sufficient to demonstrate key pipeline stages: fetch, register renaming, out-of-order dispatch, execution, completion, commit, speculative branching, and memory access. FREESS includes three step-by-step, illustrated examples that visually demonstrate how multiple instructions can be issued and executed in parallel within a single cycle. Being open source, FREESS encourages students and educators to experiment freely by writing and analyzing their own instruction-level programs and superscalar architectures.

CCS Concepts

- **Computer systems organization** → **Superscalar architectures**; *Reduced instruction set computing*;
- **General and reference** → *Evaluation*;
- **Applied computing** → **Computer-assisted instruction**.

Keywords

Superscalar, Simulator, RISC-V, Tomasulo

ACM Reference Format:

Roberto Giorgi. 2025. FREESS: An Educational Simulator of a RISC-V-Inspired Superscalar Processor Based on Tomasulo’s Algorithm. In *Proceedings of Workshop on Computer Architecture Education (WCAE ’25)*. ACM, New York, NY, USA, 8 pages. <https://doi.org/XXXXXXX.XXXXXXX>

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

WCAE '25, Tokyo, Japan

© 2025 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 978-1-4503-XXXX-X/2025/06

<https://doi.org/XXXXXXX.XXXXXXX>

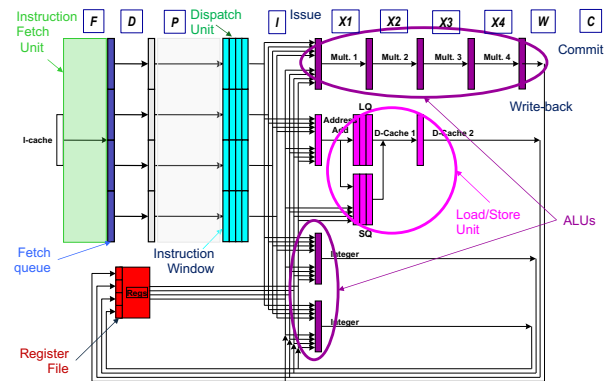


Figure 1: Structure of an Out-Of-Order Processor in FREES.

1 Introduction

Most modern microprocessors in medium- and high-end systems adopt *superscalar* architectures, building upon dynamic scheduling mechanisms such as *Thornton's scoreboard* [10] and *Tomasulo's algorithm* [11]. In this paper, we focus on Tomasulo's to illustrate how machine instructions can be executed *in parallel*, transparently to the user. These techniques are commonly known as *dynamic scheduling*, *out-of-order execution*, or *restricted dataflow* [4].

This topic is therefore a cornerstone for *Advanced Computer Architecture* courses, where *Instruction-Level Parallelism (ILP)* plays a key role in achieving *High-Performance Computing* and in accelerating *single-threaded* execution.

Superscalar architectures provide an elegant hardware-based solution to the problem of tracking control and data dependencies among instructions. While thread- and data-level parallelism are useful, executing multiple instructions per cycle can also significantly boost performance. Normally, one instruction takes several cycles to complete. A first performance improvement comes from *pipelining*, but a single pipeline can still issue at most one instruction per cycle, also known as *Flynn's bottleneck*. This limitation is overcome by issuing multiple instructions to independent functional units. This goal is achieved in *superscalar processors* (via hardware scheduling) or *VLIW architectures* (via software scheduling).

The superscalar approach based on Tomasulo’s extended algorithm augments the pipeline with several hardware structures that dynamically track data dependencies and execute instructions according to the *dataflow* principle: firing them as soon as their operands are ready. The most relevant structures involved in the tracking dependencies include: *Physical Registers (Px)*, the *Free Pool (FP)*, the *Register Map (RM)*, the *Instruction Window (IW)*, the *Reorder Buffer (ROB)*, and the *Load/Store Queues (LSQs)* (see Fig. 1).

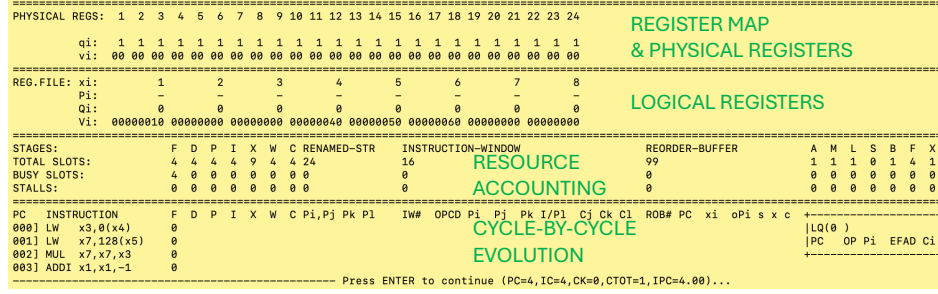


Figure 2: Cycle-by-cycle simulation progress: the screen shows the state of the superscalar’s most relevant internal structures.

Several superscalar simulators have been proposed to aid in teaching these concepts. These simulators effectively show the internal behavior of these hardware components. However, the information presented is often difficult to reproduce with pencil and paper, as updates are either overwritten on screen or dispersed across multiple disconnected windows.

FREES (Free Educational Superscalar Simulator) addresses these limitations by providing a *single screen* visualization summarizing both the current state and the cycle-by-cycle evolution of instructions through the superscalar pipeline. This evolution matches exactly what can be written on a single sheet of paper (see Fig. 2) while solving one of the exercises used for student training and for preparation verification (exams). Moreover, the output of FREES can be printed on paper for reviewing or studying purposes, even without the need to interact with a computer. This approach encourages students to bridge theoretical concepts with manual tracing and reinforces the broader course’s emphasis on performance-critical architectural design.

The syntax of the code recalls RISC-V instructions and registers due to the popularity of this novel instruction set, which is not bound to a single manufacturer and is widely adopted in Computer Architecture classes. FREES supports a small but sufficient set of seven instructions—ADD, ADDI, BEQ, BNE, LW, MUL, and SW—following a minimalistic approach in the spirit of other well-known educational tools like the LC3 simulator [7]. Although adding more instructions is very simple in FREES, the target of the simulator is studying the superscalar internals, not exploring a wide range of instructions, substitute debuggers, or other types of simulators.

Contributions

The main contributions of this work are:

- to illustrate a methodology for teaching superscalar execution, where students can manually trace execution on paper using a layout that mirrors the one provided by the tool;
- to present FREES, a simulator that provides a cycle-accurate view of the key hardware structures used in a Tomasulo-based superscalar processor;
- to support the open source release of FREES, along with guidance on writing and visualizing RISC-V-like programs.

The rest of the paper is organized as follows: Section 2 discusses related work; Section 3 describes the simulator; Section 4 presents illustrative examples; Section 5 analyzes the educational impact; and Section 6 concludes the paper.

2 Related Work

Several educational simulators have been developed to support the teaching of superscalar architecture concepts, each offering various degrees of visualization and interactivity.

SIMDE [1] is designed to support dynamic and static scheduling through Tomasulo’s algorithm and scoreboarding. It enables students to explore the flow of instructions and the hazard resolution mechanisms, providing insight into out-of-order execution and register renaming. Although SIMDE effectively represents data hazards and scheduling logic, its visualization primarily focuses on instruction status tables and resource usage. It lacks a complete, unified cycle-by-cycle overview of the instruction pipeline.

SATSim [12] provides an interactive, GUI-based environment to understand superscalar architectures. It emphasizes visual tracing of individual instructions and the state of internal buffers, which is valuable for observing execution behavior. However, SATSim offers limited feedback on stall causes or pipeline-wide performance trends, and it does not model the Load and Store Queues.

PSATSim [9] builds on SATSim by incorporating power and performance metrics. It allows users to explore the effects of microarchitectural changes on energy efficiency and throughput. Despite its improvements, PSATSim lacks comprehensive per-cycle visualization of pipeline stages.

Jaros [5] introduces a Web-based RISC-V simulator with superscalar support. Although accessible and platform-independent, this simulator primarily focuses on instruction execution and register state. Its support for architectural configuration is limited and does not include detailed memory pipeline modeling or stall diagnostics.

Other simulators like Ripes [8] and several others [6] model RISC-V pipeline, but lack a detailed superscalar modeling.

In contrast, FREES offers a more holistic approach by providing a cycle-by-cycle visualization of the full instruction pipeline, including the Instruction Window, Register Map, Reorder Buffer, and particularly the Load and Store Queues, which are often omitted in comparable tools. FREES highlights instruction-level parallelism by showing how multiple instructions proceed concurrently through fetch, rename, issue, execute, complete, and commit stages. Crucially, it also reports stall conditions and their causes, such as structural, data, or control hazards, making it easier for students to understand pipeline bottlenecks and dependency resolution. This level of feedback helps bridge the gap between theoretical understanding and practical insight into superscalar processor behavior.

Consider the following snippet of code running on 4-ways out-of-order superscalar processor. Initially, $x1=0x0010$, $x4=0x0040$, $x5=0x0050$, $x6=0x0060$ and the other registers contain zero.

```

lab1: LW  x3,0(x4)
      LW  x7,128(x5)
      MUL x7,x7,x3
      ADDI x1,x1,-1
      SW  x7,256(x6)
      ADDI x2,x2,8
      BNE x1,x0,lab1

```

Working hypothesis:

- * the fetch, decode and commit stages are 4 instructions wide
- * the instruction window has 16 slots
- * we have 24 physical registers in the free pool (excluding P0 which is hardwired to 0)
- * the reorder buffer has unlimited size
- * the integer multiplier has 4 stages
- * the load/store queues have 3 slots each and a common effective-address calculation unit
- * there are 1 ALUs for arithmetic and logic operations and for branching
- * an ALU/BANCH performs its operation in the same cycle when the operation is issued
- * reads require 1 clock cycle (after the addressing phase)
- * the register file has 4 input- and 4 output-ports
- * there are 8 logical registers (excluding x0 which is hardwired to 0)
- * the store operation leaves the issue stage as it is inserted in the store queue
- * when the X stage finds a free slot in the LQ/SQ, also a slot in the I stage is reserved
- * branches are predicted taken

Calculate the total cycles, needed to execute 3 iterations of the above loop on such machine; complete the following chart until the end of the third iteration of the code fragment above, including the renamed stream the precise evolution of the free pool of the physical registers (the register map), the Instruction Window, the Reorder Buffer (ROB) and the Load Queue (LQ) and Store Queue (SQ).

Note: Ci, Cj, Ck, Cl indicate the cycle when the corresponding physical registers are available in IW

Figure 3: At start, FREES generates the text of an exercise with specific hypotheses on the architecture of the superscalar. For example, here a 16-slot instruction window, 24 physical registers and other key parameters have been specified.

3 Description

3.1 Launching the simulator

FREES is written in pure C code, compatible with GCC from version 2.7 (year 2009) to present (year 2025) without a single warning. This characteristic makes the tool available on many computing platforms, including MS-Windows via the Ubuntu shell, for example, and, of course, all Linux-based systems. The tool is launched on the command line, and it first generates the text of the exercise, indicating the working hypothesis, based on the default options or the optional command parameters (Fig. 3). This feature is also useful to teachers for generating new exercises and for students to explore and recall the simulated architecture's key parameters. In Fig. 4 and Fig. 5, some architectural parameters are shown on architectural sketch and listed with a more detailed description.




Figure 4: FREES Parameters.

FREES PARAMETERS	UNIT	MEANING
-fw number	<inst/cycle>	Number of instruction that can be fetched in a cycle
-dw number	<insts/cycle>	Number of instruction that can be decoded in a cycle
-lw number	<insts/cycle>	Number of instruction that can be issued in a cycle
-cw number	<inst/cycle>	Number of instruction that can be committed in a cycle
-wins size	<insts>	Number of slots in the Instruction Window (IW)
-rbs size	<inst>	Number of slots in the Re-Order Buffer (ROB)
-lreg number	<num>	Number of Logical Registers
-preg number	<num>	Number of Physical Registers
-lq size	<insts>	Number of slots in the load queue (LQ)
-lfu number	<num>	Number of load functional units
-llat latency	<cycle>	Number of cycles to perform a load operation
-sq size	<insts>	Number of slots in the store queue (SQ)
-sfu number	<num>	Number of store functional units
-slat latency	<cycle>	Number of cycles to perform a store operation
-afu number	<num>	Number of integer ALUs
-alat latency	<cycle>	Number of cycles to perform an ALU operation
-mfu number	<num>	Number of integer Multipliers
-mlat latency	<cycle>	Number of cycles to perform a integer Multiplication operation

Figure 5: In yellow, the most important simulator parameters (many more parameters are available).

In Fig. 2 from top to down, the state of the superscalar parameters can be located and monitored cycle-by-cycle: the first group indicates the register map (a star is appearing below the physical registers that are allocated), qi indicates if the register is free, and vi is the register content. The second group indicates the logical registers (xi), the associated physical register (Pi), whether the value is actual (Qi), and its value (Vi).

The third group represents the accounting of the resources in terms of slots in the buffer stages (F, D, P, I, X, W, C)¹, in the renaming logic, in the instruction window, in the reorder buffer, in each of the functional units (A, M, L, S, B, F, X)². The fourth group presents for each instruction: the dynamic program counter (PC), the cycle when the instruction enters a certain stage, and other info detailed in the paper's text and the next figures.

¹The letters indicate the classic stages: F=Fetch, D=Decode/Renaming, P=Dispatch, I=Issue, X=Execute, W=Write-back and C=Commit.

²The letters indicate the functional units: A=ALU, M=integer-multiplier, L=Load, S=Store, B=branch, F=floating-point add/sub, X=floating-point-mul/div.

Table 1: FREESS instructions

Mnemonic	Operation Code (opcode)
ADD	1
ADDI	2
LW	3
SW	4
BEQ	5
BNE	6
MUL	7

3.2 Defining machine-code programs

Currently, the simulator deliberately avoids the complexity of an assembly parser. The student can code a few assembly instructions manually instead. It is left to the open source community to pick up on this point and eventually add a more complete parser. Table 1 reports the opcodes for the seven supported instructions.

To write a program, students simply replace each mnemonic with its corresponding opcode, specify the register indices, and provide any required immediate values. The handling of branch instructions is particularly instructive: instead of using labels as in assembly, students must enter the immediate value representing the number of instructions to jump—positive for forward branches, negative for backward ones, effectively replacing the role of labels in BEQ and BNE instructions.

```

loop: x3    <- mem(0+x4)    # load b(i)
      x7    <- mem(128+x5)  # load c(i)
      x7    <- x7 * x3      # b(i) * c(i)
      x1    <- x1 - 1       # decre. counter
      mem(256+x6) <- x7     # store a(i)
      x2    <- x2 + 8       # bump index
      P     <- loop; x1!=0  # close loop

```

Figure 6: Example-1. A simple loop that adds two vectors element by element and stores the results in a third vector.

Assuming the code of Fig. 6 the resulting code is the following:

```

1 3 4 0
1 7 5 128
7 7 7 3
6 1 1 -1
2 7 6 256
6 2 2 8
4 1 0 -7

```

As the processor can only execute machine code, one more bonus of FREESS is that it forces the student to realize that (or refreshing the concept if learned in a pre-requisite course) by manually converting the program from assembly mnemonic to a specific (simple) machine code format.


Superscalar execution can be analyzed in depth with assembly programs of just a few instructions (e.g., 5 to 10 instructions) and in a loop to see the effect of the branch speculation. Therefore, it is not difficult to code the program. As shown in the next Section, some pre-built examples are provided to simplify this task.

3.3 Dynamic stream and status of superscalar hardware structures

By default, the simulator assumes that the program executes three iterations of a loop (a command line parameter can change the number of iterations) and that the branch is speculatively assumed to be taken (a detailed branch predictor is future work).

3.3.1 Fetch stage. In Fig. 2, the output after the first cycle (cycle 0) is shown. The four vertical zeros under the 'F' indicate that the first four instructions are in the fetch stage. The simulator executes the next cycle every time the Enter key is pressed.

3.3.2 Decode/Rename stage. The next cycle (cycle 1) is shown in Fig. 7. The next four instructions should go to the fetch stage. However, the branch forces the fetch stage to break fetching and get only three instructions instead of four.

**Figure 7: Rename stage. The renamed instruction stream uses the newly allocated physical registers (Px).**

The first four instructions progress and go to the decode/rename stage. The columns 'Pi Pj Pk Pl' report the renamed stream and show how the logical registers are renamed to physical registers. The destination register is taken from the free pool (if not available, there will be a structural hazard), and the Register Map is updated accordingly, instruction after instruction, but within the same cycle. The Register Map (RM) tracks the assignment of the source logical to physical registers and is always visible at the top of the screen.

Please note that the screen flows: the current screen presents a complete overview of the stream's evolution, and the previous screens are still available for double-checking and can be dumped into a file for further examination and study.

3.3.3 Dispatch stage. The dispatch stage receives instructions from the renaming stage. Depending on the implementation, it stores them into the structures that will hold them until they are ready to execute, i.e., the Instruction Window (IW) or Reservation Stations (RS). We call them just *IW-SLOTS*. At the same time, a ROB entry is allocated. If either IW or ROB is full, we have a structural hazard. Again, the number of times this happens is annotated in the accounting area and explained at the bottom of the screen.

The IW-SLOT is a record that contains the entry identifier (IW#), the opcode (OPCD), the destination register (Pi), the first source register (Pj), the second source register (Pk), and the immediate (I). To track the availability of the source register values, the "flags" Cj and Ck are used: here we extend their meaning: a positive number indicates the cycle when the corresponding physical register (j or k) received its value in the IW³.

³A '-' means that the value is yet to be produced, so the instruction cannot be issued.

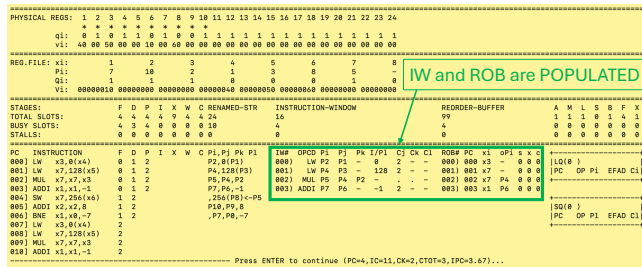


Figure 8: The Instruction Window and ROB entries are allocated in the dispatch stage.

In the example of Fig. 8 instructions 0, 1, and 2 (first LW, second LW and ADDI) received their source value at cycle 2, i.e., when they entered the IW⁴.

When an IW-slot is allocated, a corresponding ROB-SLOT is. This can be easily identified on the screen since the IW-SLOT and the ROB-SLOT are on the same row, one beside the other.

The corresponding ROB-SLOT is a record that contains the entry identifier (ROB#), the PC of the instruction, the destination logical register (xi) and eventually the old physical register (oPi) associated with that register in the Register Map (RM): this is useful when eventually there is an exception or miss-speculation and the execution has to be rolled-back. There are also three flags: 's' to record whether the operation was associated with a store, 'x' to signal an exception, and 'c' to indicate when the instruction has completed, i.e., it has written back its result. Fig. 8 shows that the IW and ROB are populated at cycle 2.

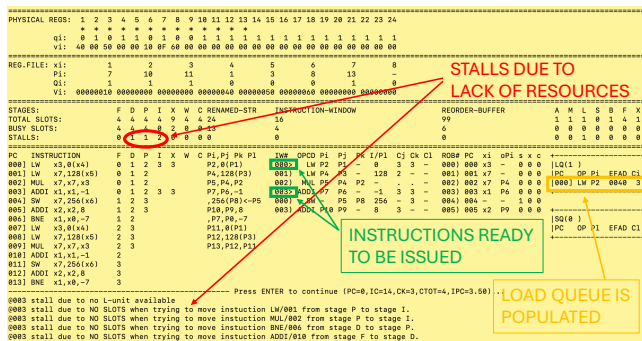


Figure 9: The third area of the screen (the accounting area) shows the stall counters to further analyze the most critical resources. Other structures are updated.

3.3.4 Issue stage. Once all source values have reached the physical registers of the instructions in the IW, one or more instructions can be issued. The issue width could be restricted to a single instruction or multiple instructions, up to the issue width (another simulation parameter). The instructions are sent to the corresponding available Functional Units (FUs) in this phase.

If the FU is occupied, we have a structural hazard, and the statistics on the screen are updated (see Fig. 9). The issued instructions

⁴We also introduced a possible third source register PI and the corresponding flag CI for future extensions.

are also marked with the character '>' besides the IW#. For example, in Fig. 9, the first LW and the ADDI are issued at cycle 3.

Load and Store Queues. Loads or stores are queued in the load queue or store queue, respectively (indicated in the right part of the screen), along their PC, opcode (OP), and the Effective Address (EFAD), which is calculated before queuing, as shown in Fig. 9. In case of a load, the value read from the memory hierarchy is written in Pi. Ci indicates the cycle when the reload is queued, and later it is updated with the cycle when the value is forwarded to the common data bus. In the case of a store, the value to be written could yet to be produced, so PI (and CI) here indicate respectively the physical register waiting for such a value and the cycle when it arrives, which corresponds to the cycle when the store is ready to access memory in the order specified by the store queue (and in synchronization with the load queue).

3.3.5 Execute stage. During the execution stage (cycle 4 of this example), we can observe that some of the instructions are fired since the IW-SLOT identifier (IW#) is becoming '-' (see Fig. 10). The related information remains on the screen for reference, while the corresponding ROB-SLOT is freed only once the instruction commits. Depending on the latency associated with the Functional Unit, the associated instruction continues to occupy the X stage for the corresponding number of cycles, assuming the unit is pipelined. In the figure Fig. 1, the multiplication consists of 4 stage (X1, X2, X3, X4) for the sake of exemplification.

For ALU operations and the EFAD calculation, the Issue (I) and Execute (X) stages often operate on the same cycle: this is our default behavior. However, as a simulation parameter, the user can also specify that I and X should always happen in different cycles.

3.3.6 Write-back (or Complete) stage. In our driving example at cycle 4, the ADDI operation completes, as can be seen, out-of-order (see Fig. 10). This is marked by c=1. However, the ROB-SLOT cannot be freed until all previous ROB-SLOTS have the c flag equal to 1, ensuring that the logical registers are updated in program order. The ROB is managed as a circular queue. In case of exceptions or mis-speculations, the operations must be undone by properly updating the Register Map and the Free Pool. Several techniques exist, but are beyond our scope here.

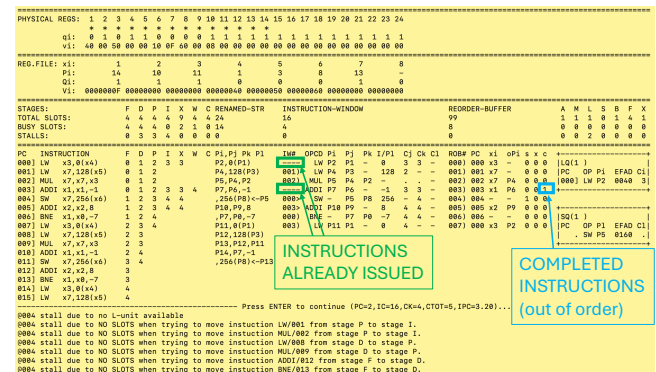


Figure 10: Annotating issued instructions in IW and completed (write-back done) are annotated in the ROB.

3.3.7 Commit stage. The student can observe how the execution progresses by observing the cycles in which each instruction enters a pipeline stage and the updates in the RM, FP, IW, ROB, LQ, and SQ. Once some instructions - up to the commit width (4 or fewer, in our case) - are completed *and* are on top of the circular queue, then those instructions are committed.

At the end of the simulation (Fig. 11), the whole evolution of the execution and the final statistics are still visible on the screen. In particular, the driving example has taken 20 cycles (CTOT) and the Instruction-Per-Cycle (IPC) is 1.05, meaning that the pipeline has been able to achieve, on average, slightly more than one instruction per cycle, despite the data dependencies and the FU latencies.

4 Examples

The FREESS package has three examples, the first described in detail in the previous section. For easier start up, the examples can be launched via a script `./run-exK.sh`, where K is '1', '2', or '3'. Once the student is more accustomed to the FREESS workflow, they can write her/his own programs (or modify the scripts), study the effect of different superscalar architectural parameters by sampling them on the command line as shown in Fig. 12.

```
$ ./freess -exe program-ex1 -pw 4 -wins 16
-pregs 24 -robs 99 -lqs 3 -sqs 3 -llat 2 -afu 1
```

Figure 12: A command line that specifies: the name of the program (-exe), the dispatch width (-pw), the window size (-wins), the number of physical registers (-pregs), the number of entries in the reorder buffer (-robs), the number of entries in the load queue (-lqs) and store queue (-sqs), the load latency (-llat), the number of arithmetic functional units (-afu).

4.1 Example-2

The second example is a smaller loop consisting of only five instructions. Fig. 13 shows the auto-generated text of this example, and the green box highlights the major modifications of the architecture: 12 IW-SLOTS, 12 Physical Registers, 12 ROB-SLOTS.

```
Consider the following snippet of code running on 4-ways out-of-order superscalar processor.
Initially, x1=0x1000, x3=0x3000, x7=0x0003 and the other registers contain zero.

lab1: LW x2,0(x1)
      ADDI x2,x2,1
      SW x2,0(x1)
      ADDI x1,x1,4
      BNE x2,x0,lab1

Working hypothesis:
* the fetch, decode and commit stages are 4 instructions wide
* the instruction window has 12 slots
* we have 12 physical registers in the free pool (excluding P0 which is hardwired to 0)
* the reorder buffer has 12 entries
* the integer multiplier has 4 stages
* the load/store queues have 3 slots each and a common effective-address calculation unit
* there are 4 ALUs for arithmetic and logic operations and for branching
* an ALU/BANCH performs its operation in the same cycle when the operation is issued
* reads require 0 clock cycle (after the addressing phase)
* the register file has 4 input- and 4 output-ports
* there are 8 logical registers (excluding x0 which is hardwired to 0)
* the store operation leaves the issue stage as it is inserted in the store queue
* when the X stage finds a free slot in the LQ/SQ, also a slot in the I stage is reserved
* branches are predicted taken

Calculate the total cycles, needed to execute 3 iterations of the above loop on such machine;
complete the following chart until the end of the third iteration of the code fragment above, including the
renamed stream the precise evolution of the free pool of the physical registers (the register map), the
Instruction Window, the Reorder Buffer (ROB) and the Load Queue (LQ) and Store Queue (SQ).
Note: C, Cj, Ck, Cl indicate the cycle when the corresponding physical registers are available in IW
```

Figure 13: Example-2. Here, a 12-slot instruction window, 12 physical registers, and 12 entries in the ROB are used.

As can be seen in the final output windows, in this case, the superscalar can achieve an IPC of 1.36 with 11 cycles of execution. This means that several more instructions are now run in parallel.

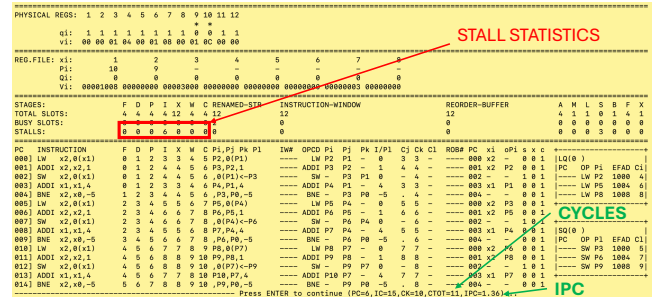


Figure 14: Example-2. Final output of the simulation of the second exercise. In this case, IPC is 1.36.

Comparing Fig. 13 and Fig. 14, we can observe that the second program generates fewer stalls due to lacking resources. In the first example (Fig. 13), we got 9 stalls due to renaming (D stage), 12 due to dispatch (P stage), 16 due to Issue (I stage), and 7 due to commit (C stage). In the second examples (Fig. 14), these numbers are respectively 0 (D), 0 (P), 6 (I), 0 (C)⁵.

4.2 Example-3

To confirm the conclusions of Example-2, another slightly different example is considered: Example-3. In this case, the program remains the same as in the Example-2, but the superscalar width is reduced to simulate a 2-way superscalar. We omit the auto-generated text to save some space, and we report the final statistics of the execution in Fig. 15. In this case, the stalls are 0(D), 3 (P), 3 (I), 0 (C). While the total number of stalls is the same (6) in example-2 and example-3, the ability to process and commit fewer instructions per cycle is limiting the IPC to 1.07 with 14 total executed cycles.




Figure 15: Example-3. For the same program as in Example-2, the IPC is now only 1.07 due to the limited width (2-way).

4.3 Analyzing the Stall Reasons

While the stall explanation is reported in the bottom part of the screen cycle-by-cycle (see Fig. 9 and Fig. 10), the same entries (all

⁵We are reporting only those statistics that are relevant for the comparison with the previous example.

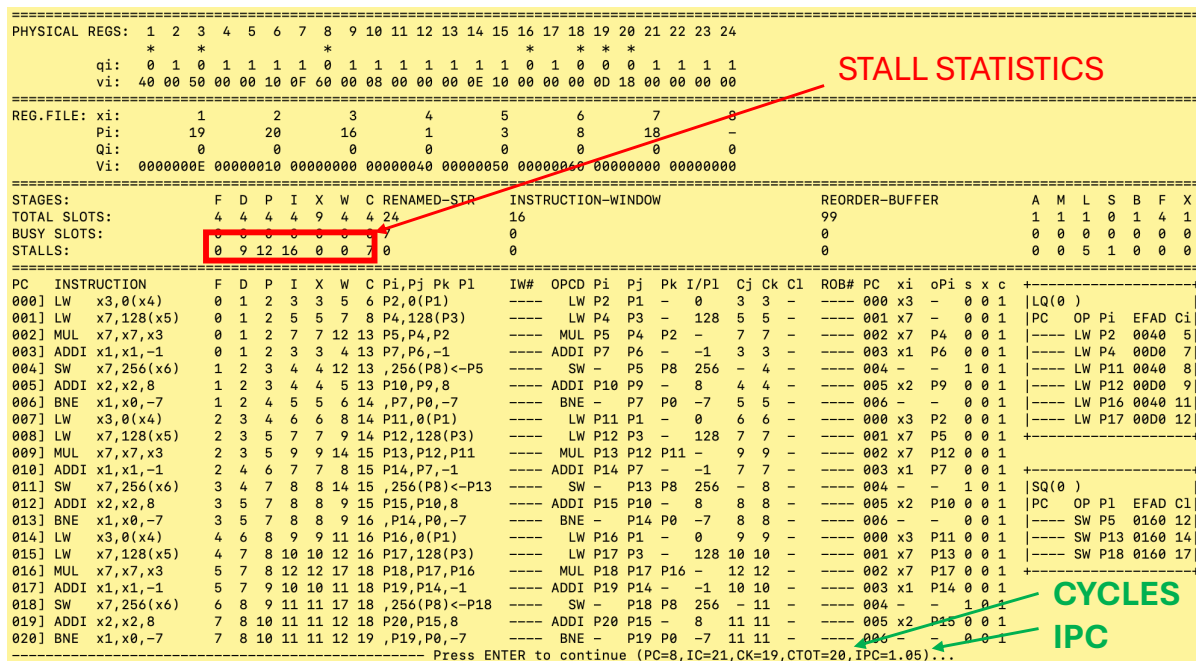


Figure 11: The final execution screen shows a summary of what happened during the program's execution. In particular, the execution of 21 instructions achieved an IPC of 1.05 (almost one instruction per cycle) in 20 cycles.

stalls) are also logged in the `stall.log` file for final reference (see Fig. 16). Each stall explanation reports: the cycle when the stall happened, the reason for the stall, the involved instruction, and the stage where this happened. By analyzing these explanations, students gain deeper insight into whether stalls are caused by resource constraints or particular instructions that could be optimized in the source program. Experimenting with architectural parameters, for example in Example-3, they can see how increasing the dispatch width, issue width, and the number of ALUs (e.g., using `-pw 3 -iw 3 -afu 2`) eliminates dispatch-stage stalls. Comparing across different configurations (such as 2-way versus 4-way superscalar) further highlights how architectural decisions directly affect performance. This exercise provides students with a hands-on understanding of how hardware limitations influence stalls and hence performance.

```

0003 stall due to NO LUnit available
0003 stall due to NO SLOTS when trying to move instruction LW/001 from stage P to stage I.
0003 stall due to NO SLOTS when trying to move instruction MUL/002 from stage P to stage I.
0003 stall due to NO SLOTS when trying to move instruction BNE/006 from stage D to stage P.
0003 stall due to NO SLOTS when trying to move instruction ADDI/010 from stage F to stage D.
0004 stall due to NO L-unit available
0004 stall due to NO SLOTS when trying to move instruction LW/001 from stage P to stage I.
0004 stall due to NO SLOTS when trying to move instruction MUL/002 from stage P to stage I.
0004 stall due to NO SLOTS when trying to move instruction LW/008 from stage D to stage P.
0004 stall due to NO SLOTS when trying to move instruction MUL/009 from stage D to stage P.
0004 stall due to NO SLOTS when trying to move instruction ADDI/012 from stage F to stage D.
0004 stall due to NO SLOTS when trying to move instruction BNE/013 from stage F to stage D.
0005 stall due to NO SLOTS when trying to move instruction ADDI/003 from stage W to stage C.
0005 stall due to NO SLOTS when trying to move instruction MUL/002 from stage P to stage I.

```

Figure 16: The log file that lists all the stall reasons that happened at a certain cycle (e.g., @003).

5 Impact

Superscalar processors and out-of-order execution are key topics in modern computer architecture courses, forming the foundation of

high-performance CPU design. These topics are commonly covered in graduate curricula using textbooks such as *Computer Architecture: A Quantitative Approach* by Hennessy and Patterson [3] and *Parallel Computer Organization and Design* by Dubois, Annavaram, Stenström [2]. However, while superscalar execution is often supported by visual tools and exercises, the more advanced concepts of superscalar execution, register renaming, and instruction reordering are harder to teach and visualize.

FRESS (Free Educational Superscalar Simulator) addresses this pedagogical gap by offering a lightweight, open-source tool to support teaching Tomasulo-style, out-of-order superscalar execution. It provides a clear, cycle-by-cycle visualization of how instructions move through the pipeline—from fetch to commit, while showing internal structures such as the Instruction Window (IW), Reorder Buffer (ROB), Register Map (RM), Free Pool (FP), and Load/Store Queues (LSQs). FRESS displays the current state of each instruction and hardware structure and highlights the causes and frequency of stalls (e.g., structural, data, and control hazards), offering a practical and detailed perspective on pipeline bottlenecks.

The RISC-V community has grown exponentially, as has the demand for more performance in RISC-V implementations for High-Performance Computing. Therefore, several RISC-V-based efforts worldwide are exploring superscalar designs to improve the execution performance. Using a RISC-V-inspired instruction set, FREESS aligns with the increasing shift in academia from older ISAs like MIPS toward the open and modern RISC-V standard. The simplified encoding of instructions and manual entry of opcode and register indices helps students understand the fundamentals of instruction encoding and control flow.

FREES has been used in the Computer Architecture curriculum at our institution since 2010 and is distributed along with illustrative examples and configuration scripts. It enables students to engage with realistic yet manageable exercises that reflect the behavior of real-world superscalar processors. The tool supports active learning, allowing students to replicate the simulation outputs on paper for deepened understanding. About 4 hours of lessons and 6 hours of practicing are planned for teaching dynamic scheduling and superscalar concept (plus 2 hours on branch prediction) at the University of Siena (course site: <https://hpca.dii.unisi.it/>). The related slides used for the teaching are available on demand.

In our experience, students initially struggle to understand out-of-order execution, but after using FREES to manually step through cycle-by-cycle outputs, they consistently report feeling more confident in tracing out-of-order execution. They appreciate how the textual interface matches what they do in paper-based exercises, and in seconds, they can evaluate the effectiveness of different architectural choices in the structure of the superscalar.

FREES also has limitations in that it does not model everything currently available in the current superscalar processor, but only the main structures typically addressed at the level of an advanced course in Computer Architecture. FREES is not a production tool, meaning there is a lot of space for improving it for different needs. The fact that it is written in pure C and has a very limited complexity of about 2000 lines of code (including lots of comments and pretty-printing functions) should make its extension easy to perform.

Finally, FREES open-source nature encourages contributions and extensions. We envision a growing community of educators and students around FREES, sharing new exercises, architectural variants, and features, further enhancing the tool's utility and educational reach.

6 Conclusions

FREES (Free Educational Superscalar Simulator) provides an effective and accessible tool for teaching the principles of superscalar processors and out-of-order execution, which are fundamental to modern computer architecture education. By offering a cycle-accurate visualization of key hardware structures such as the Instruction Window (IW), Reorder Buffer (ROB), Register Map (RM), Free Pool (FP), and Load/Store Queues (LSQs), FREES bridges the gap between theoretical concepts and practical understanding. Its unified, text-based interface allows students to manually trace execution steps, reinforcing their comprehension of dynamic scheduling and dependency resolution.

The simulator's minimalistic RISC-V-inspired instruction set simplifies learning while maintaining relevance to contemporary architectures. FREES ability to dynamically configure architectural parameters and log stall conditions enables students to explore the impact of resource limitations and pipeline hazards, fostering deeper insights into performance bottlenecks. Including illustrative examples and open-source availability further enhances its utility as an educational resource.

By abstracting away unnecessary complexity and focusing on a minimal but representative instruction set, FREES lowers the barrier for understanding key pipeline stages while maintaining technical accuracy. Its visual and interactive approach encourages

engagement, experimentation, and a deeper conceptual grasp of modern processor design. Its lightweight implementation in pure C ensures broad compatibility and ease of extension, inviting contributions from the educational community.

Avoiding the use of a GUI has the further advantage of simply printing the output of the whole evolution of the execution on a cycle-by-cycle basis or an interesting part of it, for off-line study.

Future work includes extending FREES with additional RISC-V instructions and developing more sophisticated branch prediction strategies. These enhancements will further enhance its value as both a teaching aid and a platform for architectural prototyping.

Another future possibility is to convert the C-code to a WebAssembly version for a Web-based execution.

In summary, FREES addresses the challenges of teaching superscalar execution by providing a hands-on, interactive, and adaptable platform. Its open-source nature and alignment with RISC-V position it as a valuable resource for educators and students, promoting active learning and experimentation in Computer Architecture.

Acknowledgments

I want to thank the anonymous reviewers for their encouraging comments and Jonnatan Mendoza of BSC for the useful feedback. The European Commission partially supported this work under the projects: AXIOM H2020 (id. 645496), TERAFLUX (id. 249013), HiPEAC (id. 101069836) and EDGE-ME under the Next Generation EU via the Italian National Recovery and Resilience Plan M4C2-Inv.1.4, CUP J33C22001170001, (CN00000013 - partnership ICSC).

References

- [1] I. Castilla, L. Moreno, C. González, J. Sigut, and E. González. 2007. SIMDE: An Educational Simulator of ILP Architectures with Dynamic and Static Scheduling. *Comp. App. in Eng. Education* 15, 4 (2007), 309–318. doi:10.1002/cae.20154
- [2] Michel Dubois, Murali Annamaram, and Per Stenström. 2012. *Parallel Computer Organization and Design*. Cambridge University Press, Cambridge.
- [3] John L. Hennessy and David A. Patterson. 2017. *Computer Architecture, Sixth Edition: A Quantitative Approach* (6th ed.). MKP Inc., San Francisco, CA, USA.
- [4] W. Hwu and Y. N. Patt. 1986. HPSm, a high performance restricted data flow architecture having minimal functionality. *SIGARCH Comput. Archit. News* 14, 2 (May 1986), 297–306. doi:10.1145/17356.17391
- [5] J. Jaros. 2024. Web-Based Simulator of Superscalar RISC-V Processors. In *ICS'24*. IEEE, Piscataway, NJ, USA, 1–6. doi:10.1109/SCW63240.2024.00209
- [6] G. Mariotti and R. Giorgi. 2022. WebRISC-V: A 32/64-bit RISC-V pipeline simulation tool. *ELSEVIER SoftwareX* 18 (May 2022), 1–7. doi:10.1016/j.softx.2022.101105
- [7] Yale Patt and Sanjay Patel. 2004. *Introduction to Computing Systems: From Bits and Gates to C and Beyond* (2 ed.). McGraw-Hill, New York, NY, USA.
- [8] Morten B. Petersen. 2021. Ripes: A Visual Computer Architecture Simulator. In *ISCA-WCAE'21*. IEEE, Virtual Conference, 1–8.
- [9] C. W. Smullen. 2006. PSATSim: An Interactive Graphical Superscalar Architecture Simulator for Power and Performance Analysis. In *ISCA-WCAE'06*. ACM, New York, NY, USA, 1–6. doi:10.1145/1275620.1275627
- [10] James E. Thornton. [n. d.]. *Design of a Computer: The Control Data 6600*. Scott, Foresman and Co., Glenview, IL, USA.
- [11] Robert M. Tomasulo. 1967. An Efficient Algorithm for Exploiting Multiple Arithmetic Units. *IBM Journal of Research and Development* 11, 1 (1967), 25–33.
- [12] S. Wolff. 2000. SATSim: A Superscalar Architecture Trace Simulator Using Interactive Visualization. In *ISCA-WCAE'00*. ACM, New York, NY, USA, 1–7. doi:10.1145/1275240.1275249

A Online Resources

The FREES Educational Simulator for RISC-V inspired Superscalar Processors based on Tomasulo's Algorithm is available at this address: <https://github.com/robgorgi/frees>

Received 19 May 2025