

CPU simulator FlyHip

Branko Šter¹, Paul Daubin²

¹University of Ljubljana, Faculty of Computer and Information Science, Večna pot 113, 1000, Ljubljana

²ESIGELEC Engineering school, Av. Galilée, 76800 Saint-Étienne-du-Rouvray, France

E-mail: branko.ster@fri.uni-lj.si

Abstract

An implementation of simulator of processor HIP, a little simplified version of MIPS I architecture, is described. Both non-pipeline and pipeline versions of HIP are supported. The simulator, written in Python, can be run either in a command-line mode or with a graphical user interface (GUI), implemented in Flutter, a modern open-source user interface software development kit. It can be run by default as a desktop application, but can potentially be adapted for running in a web browser or even as a mobile application.

1 Introduction

Although x86 instruction set architectures (ISA) are widespread commercially, CISC architectures are quite complex and as such not appropriate for undergraduate level courses, therefore courses usually employ various RISC architectures. The instruction set MIPS I by MIPS Technologies is a very popular RISC ISA, frequently used in computer architecture courses all over the world. For example, the MIPS I ISA was used in processors MIPS R2000 and R3000. For the purpose of teaching computer architecture in Faculty of computer and information science at University of Ljubljana, a little simplified version of MIPS I called HIP was proposed in the book [1] and used throughout the years.

The previous simulator for HIP CPU [4] was written with GUI that was based on WinMIPS64 software [5]. It was used in courses such as Computer systems architecture and similar courses at Faculty of Computer and Information Science. However, during the pandemic we needed to program Moodle quizzes for the evaluation of the students' work. These were written in Python programming language. Inspired by these scripts, the idea for writing a new simulator emerged.

The new simulator of the processor HIP is implemented in Python, a very popular programming language for some time already. The command-line version is called pyHip (as one would expect). Since the frontend pyHip (written in Python) is connected with backend Flutter GUI, we named the complete program FlyHip (Flutter + pyHip).

The reason for developing the new simulator was not only the unavailability of the source code of the old one,

but rather a greater flexibility of maintaining and upgrading one's own software. Additionally, since Flutter is a modern and very popular multi-platform open-source graphical user interface (GUI) software development kit (by Google), the GUI is way more flexible and expandable than the Windows GUI of WinMIPS64, which had to be frequently running in some 'compatibility mode' for older Windows versions, especially due to not flexible font size. In Flutter, more features and window-related options are available.

2 HIP CPU

HIP instruction set architecture is a little simplified MIPS I ISA, as such being a typical RISC CPU. It has a load/store architecture, meaning that ALU operations directly on operands in memory are not possible, so the operands first have to be loaded from the memory to registers in order to perform operations on them.

HIP supports only two instruction formats (Fig. 1). HIP ISA has 52 instructions: 31 in format 1 and 21 in format 2. Format 1 has one source register (Rs1), one input immediate operand (a constant), and a destination register (Rd). This format is used in all load/store instructions, as well as in all ALU instructions that take one immediate operand. In the load/store instructions, the immediate operand serves as the offset added to the base register (base addressing). Format 2 is used in all ALU instructions that take two input operands in registers.




Figure 1: HIP supports two instruction formats: format 1 .

The HIP instructions may be divided into 4 groups:

- Load/store instructions
- ALU instructions:
 - Arithmetic operations add and subtract

- Logical bitwise operations
- Shift operations
- Compare-and-Set operations

- Control instructions
- System instructions

All the instructions are listed in Table 1. There are no

Table 1: HIP instruction set architecture.

Load/store instructions: LB (load byte), LH (load halfword), LW (load word), LBU (LB unsigned), LHU (LH uns.), SB (store byte), SH (store halfword), SW (store word)
ALU instructions: <i>Arithmetic operations add and subtract:</i> ADD (add signed), SUB (subtract sig.), ADDU (add uns.), SUBU, ADDI (add signed immediate), SUBI (sub. signed imm.), ADDUI (add uns. imm.), SUBUI (sub. uns. imm.) <i>Logical bitwise operations:</i> AND, OR, XOR, NOT (one's complement), ANDI (imm.), ORI, XORI <i>Shift operations:</i> SLL (shift left logical), SRL (sh. right log.), SRA (sh. right arith.), SLLI (SLL imm.), SRLI (SRL imm.), SRAI, LHI (load high imm.) <i>Compare-and-Set operations:</i> SEQ (set if equal), SNE (set if not eq.), SLT (set if less than), SGT (set if greater than), SLTU (SLT uns.), SGTU (SGT uns.), SEQI (SEQ imm.), SNEI, SLTI, SGTI, SLTUI, SGTUI
Control instructions: BNE (branch if not equal zero), BEQ (branch if eq. zero), J (jump, unconditional), CALL (jump to subroutine), TRAP (jump to vector address), RFE (return from exception)
System instructions: EI (enable interrupts), DI (disable int.), MOVER (move from EPC to register), MOVRE (move from reg. to EPC)

multiply/divide instructions and no support for floating point instructions.

The organization of the basic (non-pipeline) version is shown in Fig. 2. The control unit (on the left) implements a state machine with 28 states. The right side shows the datapath, including ALU and various registers: 32 general-purpose registers R0, R1, ..., R31, with interface registers A, B and C, and special-purpose registers, such as PC (Program Counter), EPC (Exception Program Counter), which is basically the back-up register of PC when performing interrupt service routines, MAR (Memory Address Register) and MDR (Memory Data Register).

To simplify the operation, the user does not have to define available code and data segments in the simulator settings, in contrast to simulator WinHIP, where students always have to make sure that these settings coincide. In pyHip (and FlyHip), the simulator automatically allocates the segments given in the assembly source. A specified default size of the segments can be set in advance.

3 Simulator

The simulator is written in Python programming language. The command-line version is called *pyHip*. It loads a user-defined HIP assembly source file, runs it through lexer and parser, which are implemented in a manner similar to [6, 7].

The non-pipeline version executes instructions sequentially one after another. Each instruction has its specific length in the number of clock cycles - caused by the specific CPU organization of HIP (Fig. 2), as defined in [1]. The instructions last form 3 to 6 clock cycles. Depending

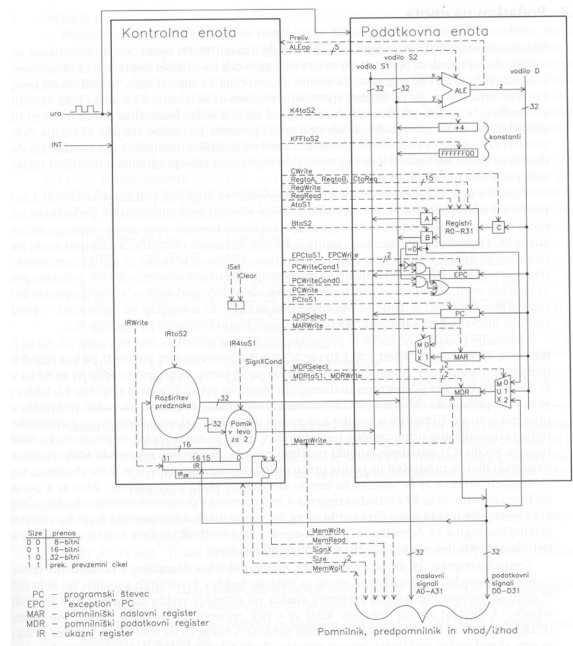


Figure 2: Non-pipeline HIP CPU (from [1]).

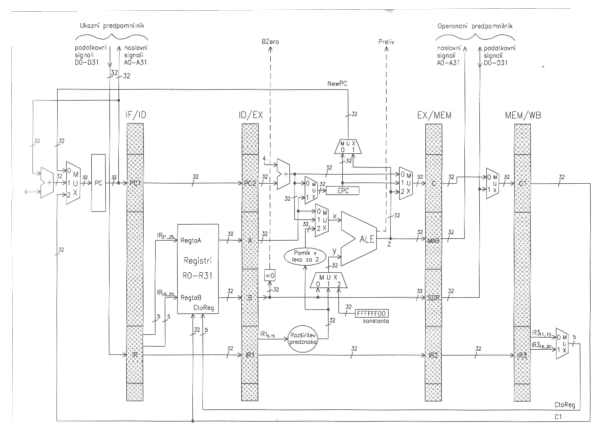


Figure 3: Pipeline HIP CPU datapath (from [1]).

on a program running on HIP, the average CPI (Clock Per Instruction) is usually between 4 and 5. Although the non-pipeline simulator does not operate on the cycle-level, it still accumulates the number of clock cycles for the statistics, such as the execution time and the average CPI of a program. An example of the HIP assembly code:

```
.data
.org 0x400
POD:  byte 9,2,0,4,7,0,1,2,0,0,0,0

.code
.org 0x0
addui r9, r0, POD
LOOP: lb r1, 0(r9)
      lb r2, 1(r9)
      sub r3, r1, r2
      subi r3, r3, 1
      sb 2(r9), r3
      slti r17, r3, 0
      addui r9, r9, 3
      beq r17, LOOP
      halt
```

Fig. 4 shows an example of the command-line pyHip output.

```
(pyhip) a
Instruction 9 :
0x1c: 0x08420000 ADDUI R2, R2, #AH
Registers:
R00: 0x00000000 R04: 0xcc04f5134 R08: 0xdeddc0f3 R12: 0xe489e0e8 R16: 0x8:
R01: 0x00000001 R05: 0xa655c27d R09: 0xacce24b9 R13: 0xf69db055 R17: 0xf:
R02: 0x20000000 R06: 0x82f53feb R10: 0x10000000 R14: 0x32bb0702 R18: 0x4:
R03: 0xe49f6927 R07: 0xcbe95d54 R11: 0x57fdac69 R15: 0x99ebdae3 R19: 0xb:
PC: 0x20 , EPC: 0x0 , I: 0
Memory:
0x00000400: 0c 0c 37 fd 00 00 00 00 32 00 00 00 32 cf b3 32 66 9b a7 ff cd 74
0x00000420: ae 01 73 3a 75 8c 02 25 9f c6 2b 12 5e f7 80 a9 b2 fd b4 76 3a
Stats:
Cycles: 44, CPI: 4.400, Loads: 4, Stores: 2, Jumps: 2, Branches not taken: 0.
```

Figure 4: An example of stepping through the code with the non-pipeline command-line simulator pyHip.

On the other hand, the pipeline version has to operate on the cycle-level. HIP has a classical 5-stage pipeline, consisting of the following stages: IF (Instruction Fetch), ID (Instruction Decode), EX (EXecute), ME (MEemory), and WB (WriteBack) (Fig. 3), as can be seen in the command-line pipeline pyHip output in Fig. 5.

Since the pipeline datapath is a synchronous digital circuit, all the registers renew their state at the same time, typically at the rising edge of the clock signal. Since in a software simulator this is not possible to do directly, the synchronicity is achieved by sequential updating of different stages from right to left (in Fig. 3), due to the fact that instructions move through the pipeline from left to right, similarly to a shift register. However, a care must be taken in cases of feedback connections, for example, register C1 is 'written back' to a destination register in the WB phase.

By default, the pipeline HIP CPU detects data hazards of type RAW (Read-After-Write). Other types of data hazards, such as Write-After-Write and Write-After-Read, cannot occur in HIP due to its simplicity. RAW hazards are detected by additional logic that compares the address (number) of the destination register in the instruction in register IR3 (just prior to the WB phase) to the addresses of the source registers in phases IR2, IR1, and IR (the next three instructions). In case of the same register being written and then read a little later, it must be assured that the WB phase of the write is done before the ID (decoding) phase of the read, so that the register is actually *read after being written* by an earlier instruction. If necessary, pipeline interlocks ('bubbles') need to be inserted to ensure the correct operation, in effect by performing the NOP operations.

The pipeline supports additional options of bypassing (data-forwarding) and delayed branches. *Bypassing* requires additional logic for detecting the RAW data hazards, as just described, to determine which of the result registers (C1, C or even the ALU result) is written back directly via additional feedback connections to the source registers A and B, in order to avoid stalls.

When conditional branches in the code are actually performed (the condition is TRUE), two next instructions, which are already in the pipeline stages IF and ID, need to be cancelled, since they must not be executed. Thus two additional stalls occur whenever a branch is done. However, the option of *delayed branches* allows for re-

ordering the code - putting at most two instructions from before the branch to after the branch, in order to avoid the two stalls.

```
(pyhip) a
Clock cycle: 20
Instruction:
0x20: 0x9c11ffe0 BEQ R17, ZANKA
Registers:
R00: 0x00000000 R04: 0xd53cf16c R08: 0x7d2bf190 R12: 0xd0c04083 R16: 0x2dad1b83 R20: 0xfc:
R01: 0x00000009 R05: 0x71c62394 R09: 0x00000400 R13: 0x61a41dcb R17: 0xaa0e4bf3 R21: 0x62:
R02: 0x00000002 R06: 0xd212343e R10: 0x40f20a10 R14: 0xc0b702cc R18: 0x8c78fe7f R22: 0x0b:
R03: 0x00000006 R07: 0xa0148c68 R11: 0x4f09f139 R15: 0xa37aa48d R19: 0x85df95c9 R23: 0x0a:
PC: 0x0000024, PC1: 0x0000020, PC2: 0x000001c, EPC: 0x0000000, I: 0
IR: 9c11ffe0, IR1: 09290003, IR2: 28710000, IR3: a1230002
A: 00000400, B: 00000400, C: 00000000, MAR: 00000400, SDR: 00000400
Memory:
0x400: 09 02 06 04 07 00 ff 02 00 00 00 00 31 69 40 18 ad cf 6b 0b 78 27 71 71 44 36 77 43
0x420: db e5 3b 1d f8 85 f2 e5 39 1f 21 7c 4a 8a 6c 86 69 7c 49 70 e0 dd 2b ee 9c fd a8 bd
Pipeline execution:
addr. instruction clk: 0 1 2 3 4 5 6 7 8 9 10 11 12 13 14 15 16 17 18 19 20
0: ADDUI R9, R0, #00 IF ID EX ME WB - - - - - - - - - - - - - - - - - -
4: LB R1, 0(R9) - IF o o o ID EX ME WB - - - - - - - - - - - - - - -
8: LB R2, 1(R9) - - - - - IF ID EX ME WB - - - - - - - - - - - - - - -
c: SUB R3, R1, R2 - - - - - - IF o o o ID EX ME WB - - - - - - - - -
10: SUBI R3, R3, 1 - - - - - - - - - IF o o o ID EX ME WB - - - - -
14: SB 2(R9), R3 - - - - - - - - - - IF o o o o ID EX ME - - - - -
18: SLTI R17, R3, 0 - - - - - - - - - - - - - - - IF ID EX - - - - -
1c: ADDUI R9, R9, 3 - - - - - - - - - - - - - - - - - IF ID - - - - -
20: BEQ R17, ZANKA - - - - - - - - - - - - - - - - - - IF - - - - -
Stats:
Cycles: 21, Instructions: 5, CPI: 4.200, Loads: 2, Stores: 1, Jumps: 0, Branches not taken: 0,
RAW stalls: 12
```

Figure 5: An example of stepping through the code with the pipeline command-line simulator pyHip. In the presented case, no bypassing (data forwarding) was used, so three RAW (Read-After-Write) stalls or 'bubbles' are inserted.

4 Graphical User Interface (GUI)

4.1 Objectives of the GUI

The main objective of the GUI is to present the information in a way that is comprehensible - accessible to users and efficient: showing clearly the crucial data without bloating and allowing interaction between the user and the simulator. For this project we also wanted to be able to deploy on multiple platforms as needed. Mainly desktop platforms: Windows, Linux, and MacOS, but also on web platforms for a more lightweight experience, and eventually a smaller version on mobile.

4.2 Technologies used

To that end, we decided developing the simulator's interface using Flutter, a framework built by Google around the Dart programming language. It is a recent technology (2017), but is backed by a powerful company and thus provides a level of future proofing, while allowing us to satisfy the specifications, being able to deploy on several platforms with the same code. Desktop support was also officially released only a few months ago, making it a perfect opportunity.

Our interface, using the MVC (model-view-controller) design pattern can be decomposed into 3 main components:

- Views, where information is displayed and interacted with by the users. In Flutter, each element, from the global scaffold of the application to the smallest text, is built using Widgets, predefined and highly customisable blocks, that can be adjusted to suit the user's needs, and put one on top of another. The resulting widget tree defines how the information is accessed, displayed, and under which form.

- The controller is what allows us to predefine HTTP requests that the user is going to send to the Python backend. For example, we have set up a request, using a specific IP and port, that requests the simulator to move forward in the simulation. By simply pressing a button, the user can send the instruction to the simulator. This is also the way we get the information, in JSON form.
- The model is where this information is stocked. We define the elements we wish to extract from the JSON file, decompose it in its various information elements, and extracts the ones we wish to use. This information is mapped into multiple classes or 'models', which allows us to manipulate it.

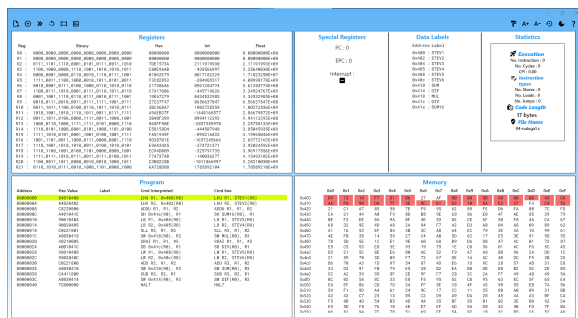


Figure 6: An example of the current GUI on desktop, showing the basic version of the simulator

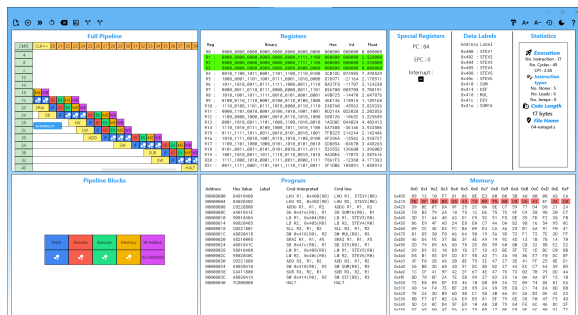


Figure 7: An example of the current GUI on desktop, showing the pipeline of the simulator (subject to change)

4.3 Features

The interface has been designed keeping both students and teachers in mind, giving users as much tools as possible but staying straightforward and clear. It is used by providing an assembly file using a file explorer and then compiled by the Python program. Then, the execution can be followed step by step, in the non-pipeline version (6) at the granularity of instructions, and in the pipeline version (7) at the granularity of clock cycles, and highlighting the important information. Alternatively, the program can be run until the end to see the end result. Then, users can either select a new file or simply rewind the current one at their convenience.

The program provides numerous features allowing the teacher to customize the desired way of showing the information to the student, including changing the colour

of the interface, zooming on important information and more. These options are intuitive, done either using the buttons on the top left, or simply scrolling or clicking on the interface. The students are able to get more information on different elements by a simple mouseover, allowing us to keep the interface clean while still keeping them afloat. Each step of the program is also highlighted to show the relevant information, so they can identify and focus on the most crucial points.

The program displays lots of information, under different forms, ranging from a simple table for registers (bottom right), to an intricate and adaptive design for the pipeline (top left), and designs in between, adapted for the data they contain. Additionally, the interface is decomposed in different screens, between which one can seamlessly navigate as in any website or application. The source code of the program is available at [8].

5 Conclusion

A CPU simulator is a very valuable tool for students learning computer architecture, since on practical examples of assembly programs students can gradually learn how the CPU behaves. They can view registers and memory, explore runtime statistics, performance issues, pipeline stalls, cache behaviour, and many more.

We describe an implementation of simulator of processor HIP, a little simplified version of processor MIPS R3000. Both non-pipeline and pipeline versions of HIP are supported. The simulator can be run either in a command-line mode or with a GUI implemented in Flutter. It can be run by default as a desktop application, but can potentially be adapted for running in a web browser, with the advantage of being more system-independent.

In the recent years, the RISC-V open standard ISA is growing in popularity, gradually also in university courses. Besides of being open-source, it is also more modern. We may use the same framework to develop also a RISC-V simulator.

6 Acknowledgements

We thank Ratko Pilipović and Nejc Ilc for their help and advices in designing the GUI.

References

- [1] D. Kodek: Arhitektura in organizacija računalniških sistemov, Bi-Tim, Šenčur, 2008.
- [2] D. A. Patterson, J. L. Hennessy: Computer organization and design: The hardware/software interface, Morgan Kaufman, 2012.
- [3] A. Akram, L. Sawalha: A Survey of Computer Architecture Simulation Techniques and Tools, IEEE Access, 2019.
- [4] D. Šonc: Simulator WinHIP, FRI, 2007-2012.
- [5] M. Scott: WinMIPS64, <http://indigo.ie/~mscott/>
- [6] A. Z. Henley: teenytinycompiler, <https://github.com/AZHenley/teentytinycompiler>
- [7] A. Z. Henley: Let's make a Teeny Tiny compiler, <https://austinhenley.com/blog/teentytinycompiler1.html>
- [8] <https://github.com/brankoster/pyhip>