





- We should use fixed instruction encoding if interested in performance, and use variable instruction encoding if interested in code size
- Should provide at least 16 general-purpose registers, be sure all addressing modes apply to all data transfer instructions and aim for a minimalist instruction set. This will reduce the complexity of writing a correct compiler (which is a major limitation on the amount of optimization that can be done).
- Before suggesting new instruction set features, look at optimized code first as a compiler might completely remove the instructions the architect tries to improve...

Example: MIPS, RISC-V follow those recommendations

- MIPS64 has 32 64-bit general purpose registers (GPR)
- Data types: 8-, 16-, 32-, and 64-bit for integers and 32-bit single precision and 64-bit double precision for floating point
- Addressing modes: immediate and displacement, both with 16-bit fields
- Types of instruction format:
 - MIPS three types: R,I, and J
 - RISC-V six types: R,I,S,SB,UJ,U
- Supports the list of simple instructions recommended plus a few others
- Control handled through a set of jumps and a set of branches



MIPS Encoding Summary Microprocessor without Interlocked Pipeline Stages • 3 instruction formats: R, I, and J types shamt funct R rs rt \mathbf{rd} op 16 bit address rt Ι op rs 26 bit address op J 6

RISC-V Encoding Summary

(Field Size) 7 b	bits	5 bits					Comments
-			5 bits	3 bits	5 bits	7 bits	
R-type fur	nct7	rs2	rs1	funct3	rd	opcode	Arithmetic instruction format
I-type	immediate[11:0] rs			funct3	rd	opcode	Loads & immediate arithmetic
S-type immed	d[11:5]	rs2	rs1	funct3	immed[4:0]	opcode	Stores
SB-type immed[[12,10:5]	rs2	rs1	funct3	immed[4:1,11]	opcode	Conditional branch format
UJ-type	immediate[20,10:1,11,19:12]				rd	opcode	Unconditional jump format
U-type	immediate[31:12]				rd	opcode	Upper immediate format





RISC-V Datapath: 5 Stages IF/ID ID/EX EX/MEM MEM/WB Shift left 1 0 M u x PC Read register 1 Read data 1 Read register 2 Registers ALU ALU Instruction Read data memory Read data 2 1 M U X 0 Addross Write register 0 M u x 1 result Data memory Write data Write data 32 lmm Gen 10

Four Compile Time Schemes to Reduce Branch Hazard Penalties due to One-Delay Stalls

- #1: Stall until branch direction is clear
- #2: Predict (or Treat) Branch as Not Taken
 - $^{\circ}\,$ Execute successor instructions in sequence as if the branch were a normal instruction
 - $^\circ$ "Squash" instructions in pipeline if branch actually taken
 - $^{\circ}$ Advantage of late pipeline state update
 - ° 47% branches not taken on average
 - $^\circ$ PC+4 already calculated, so use it to get next instruction

#3: Predict Branch as Taken

- $^\circ$ 53% branches taken on average
- ° But have not calculated branch target address
 - So, still incurs 1 cycle branch penalty

Four Compile Time Schemes to Reduce Branch Hazard Penalties due to One-Delay Stalls #4: Delayed Branch Introduce the sequential successor instruction, which is executed irrespective of whether or not the branch is taken The job of the Compiler is to make the successor instructions valid and useful

