
System I

Instruction Set Architecture

Bo Feng, Lei Wu, Rui Chang

Zhejiang University

Disclaimer

- **Many images and resources used in this lecture are collected from the Internet, and they are used only for the educational purpose. The copyright belong to the original owners, respectively.**

- **Part of slides credit to**
 - **David A. Patterson and John L. Hennessy. Computer Organization and Design RISC-V Edition: The Hardware Software Interface, 1st Edition.**
 - **John L. Hennessy and David A. Patterson. Computer Architecture: A Quantitative Approach, 6th Edition.**
 - **Andrew Waterman and David A. Patterson. The RISC-V Reader: An Open Architecture Atlas.**
 - **CSCE 513, Prof. Yonghong Yan @ University of North Carolina at Charlotte**
 - **CENG3420, Bei Yu @ The Chinese University of Hong Kong**
 - **CS 3410, Prof. Hakim Weatherspoon @ Cornell University**
 - **CS 162, Prof. Sam Kumar @ UC Berkeley**
 - **CSc 453, Prof. Saumya Debray @ University of Arizona**

Overview

- RISC-V ISA
- RISC-V Assembly Language

Overview

- RISC-V ISA
- RISC-V Assembly Language

What is RISC-V?

- RISC-V (pronounced "risk-five") is an ISA standard
 - An open-source implementation of a reduced instruction set computing (RISC) based instruction set architecture (ISA)
 - There was RISC-I, II, III, IV before
- Most ISAs: X86, ARM, Power, MIPS, SPARC
 - Commercially protected by patents
 - Preventing practical efforts to reproduce the computer systems.
- RISC-V is open
 - Permitting any person or group to construct compatible computers
 - Use associated software
- Originated in 2010 by researchers at UC Berkeley
 - Krste Asanović, David Patterson and students
- [ISA Specifications](#)
 - Unprivileged specification version 20191213 (v2.2)
 - Privileged specification version 20211203 (v1.11)
 - More on github: <https://github.com/riscv/riscv-isa-manual>



<https://riscv.org/>

Goals in Defining RISC-V

- A completely open ISA that is freely available to academia and industry
- A real ISA suitable for direct native hardware implementation, not just simulation nor binary translation
- An ISA that avoids “over-architecting” for
 - A particular microarchitecture style (e.g., microcoded, in-order, decoupled, out-of-order) or
 - Implementation technology (e.g., full-custom, ASIC, FPGA), but which allows efficient implementation in any of these
- RISC-V ISA includes
 - A small base integer ISA, usable by itself as a base for customized accelerators or for educational purposes, and
 - Optional standard extensions, to support general-purpose software development
 - Optional customer extensions
- Support for the revised 2008 IEEE-754 floating-point standard

RISC-V Principles

- Generally kept very simple and extendable
 - Whether short, long, or variable
- Separated into multiple specifications
 - User-level ISA spec (compute instructions)
 - Compressed ISA spec (16-bit instructions)
 - Privileged ISA spec (supervisor-mode instructions)
 - More...
- ISA support is given by RV + word-width + extensions supported
 - E.g., RV32I means 32-bit RISC-V with support for the I (integer) instruction set

User-Level ISA

- Defines the normal instructions needed for computation
 - A mandatory **base integer ISA**
 - **I: Integer instructions:**
 - ALU
 - Branches/jumps
 - Loads/stores
 - **Standard extensions**
 - M: Integer Multiplication and Division
 - A: Atomic Instructions
 - F: Single-Precision Floating-Point
 - D: Double-Precision Floating-Point
 - C: Compressed Instructions (16 bit)
 - **G = IMAFD: integer base + four standard extensions**
 - Optional extensions


Basic RISC-V ISA

- Both 32-bit and 64-bit address space variants
 - RV32 and RV64
- Easy to subset/extend for education/research
 - RV32IM, RV32IMA, RV32IMAFD, RV32G
- SPEC on the website
 - www.riscv.org

Name of base or extension	Functionality
RV32I	Base 32-bit integer instruction set with 32 registers
RV32E	Base 32-bit instruction set but with only 16 registers; intended for very low-end embedded applications
RV64I	Base 64-bit instruction set; all registers are 64-bits, and instructions to move 64-bit from/to the registers (LD and SD) are added
M	Adds integer multiply and divide instructions
A	Adds atomic instructions needed for concurrent processing; see Chapter 5
F	Adds single precision (32-bit) IEEE floating point, includes 32 32-bit floating point registers, instructions to load and store those registers and operate on them
D	Extends floating point to double precision, 64-bit, making the registers 64-bits, adding instructions to load, store, and operate on the registers
Q	Further extends floating point to add support for quad precision, adding 128-bit operations
L	Adds support for 64- and 128-bit decimal floating point for the IEEE standard
C	Defines a compressed version of the instruction set intended for small-memory-sized embedded applications. Defines 16-bit versions of common RV32I instructions
V	A future extension to support vector operations (see Chapter 4)
B	A future extension to support operations on bit fields
T	A future extension to support transactional memory
P	An extension to support packed SIMD instructions: see Chapter 4
RV128I	A future base instruction set providing a 128-bit address space

RISC-V Processor State

- Program counter (**PC**)
- 32 32/64-bit integer registers (**x0-x31**)
 - x0 always contains a 0
 - x1 to hold the return address on a call.
- 32 floating-point (FP) registers (**f0-f31**)
 - Each can contain a single- or double-precision FP value (32-bit or 64-bit IEEE FP)
- FP status register (**fsr**), used for FP rounding mode & exception reporting



RV32I

Integer Computation

add { immediate }

subtract

{ and
or
exclusive or } { immediate }

{ shift left logical
shift right arithmetic
shift right logical } { immediate }

load upper immediate
add upper immediate to pc

set less than { immediate } { unsigned }

Control transfer

branch { equal
not equal }

branch { greater than or equal
less than } { unsigned }

jump and link { register }

Loads and Stores

{ load
store } { byte
halfword
word }

load { byte
halfword } unsigned

Miscellaneous instructions

fence loads & stores
fence.instruction & data

environment { break
call }

control status register { read & clear bit
read & set bit
read & write } { immediate }

ALU Instructions

Example instruction	Instruction name	Meaning
add x1,x2,x3	Add	$\text{Regs}[x1] \leftarrow \text{Regs}[x2] + \text{Regs}[x3]$
addi x1,x2,3	Add immediate unsigned	$\text{Regs}[x1] \leftarrow \text{Regs}[x2] + 3$
lui x1,42	Load upper immediate	$\text{Regs}[x1] \leftarrow 0^{32} \# \# 42 \# \# 0^{12}$
sll x1,x2,5	Shift left logical	$\text{Regs}[x1] \leftarrow \text{Regs}[x2] \ll 5$
slt x1,x2,x3	Set less than	$\text{if } (\text{Regs}[x2] < \text{Regs}[x3])$ $\text{Regs}[x1] \leftarrow 1 \text{ else } \text{Regs}[x1] \leftarrow 0$

Figure A.26 The basic ALU instructions in RISC-V are available both with register-register operands and with one immediate operand. LUI uses the U-format that employs the rs1 field as part of the immediate, yielding a 20-bit immediate.

Load/Store Instructions

Example instruction	Instruction name	Meaning
ld x1, 80(x2)	Load doubleword	$\text{Regs}[x1] \leftarrow \text{Mem}[80 + \text{Regs}[x2]]$
lw x1, 60(x2)	Load word	$\text{Regs}[x1] \leftarrow_{64} \text{Mem}[60 + \text{Regs}[x2]]_0^{32} \text{##}$ $\text{Mem}[60 + \text{Regs}[x2]]$
lwu x1, 60(x2)	Load word unsigned	$\text{Regs}[x1] \leftarrow_{64} 0^{32} \text{##} \text{Mem}[60 + \text{Regs}[x2]]$
lb x1, 40(x3)	Load byte	$\text{Regs}[x1] \leftarrow_{64} (\text{Mem}[40 + \text{Regs}[x3]]_0)^{56} \text{##}$ $\text{Mem}[40 + \text{Regs}[x3]]$
lbu x1, 40(x3)	Load byte unsigned	$\text{Regs}[x1] \leftarrow_{64} 0^{56} \text{##} \text{Mem}[40 + \text{Regs}[x3]]$
lh x1, 40(x3)	Load half word	$\text{Regs}[x1] \leftarrow_{64} (\text{Mem}[40 + \text{Regs}[x3]]_0)^{48} \text{##}$ $\text{Mem}[40 + \text{Regs}[x3]]$
flw f0, 50(x3)	Load FP single	$\text{Regs}[f0] \leftarrow_{64} \text{Mem}[50 + \text{Regs}[x3]] \text{##} 0^{32}$
fld f0, 50(x2)	Load FP double	$\text{Regs}[f0] \leftarrow_{64} \text{Mem}[50 + \text{Regs}[x2]]$
sd x2, 400(x3)	Store double	$\text{Mem}[400 + \text{Regs}[x3]] \leftarrow_{64} \text{Regs}[x2]$
sw x3, 500(x4)	Store word	$\text{Mem}[500 + \text{Regs}[x4]] \leftarrow_{32} \text{Regs}[x3]_{32..63}$
fsw f0, 40(x3)	Store FP single	$\text{Mem}[40 + \text{Regs}[x3]] \leftarrow_{32} \text{Regs}[f0]_{0..31}$
fsd f0, 40(x3)	Store FP double	$\text{Mem}[40 + \text{Regs}[x3]] \leftarrow_{64} \text{Regs}[f0]$
sh x3, 502(x2)	Store half	$\text{Mem}[502 + \text{Regs}[x2]] \leftarrow_{16} \text{Regs}[x3]_{48..63}$
sb x2, 41(x3)	Store byte	$\text{Mem}[41 + \text{Regs}[x3]] \leftarrow_8 \text{Regs}[x2]_{56..63}$

Figure A.25 The load and store instructions in RISC-V. Loads shorter than 64 bits are available in both sign-extended and zero-extended forms. All memory references use a single addressing mode. Of course, both loads and stores are available for all the data types shown. Because RV64G supports double precision floating point, all single precision floating point loads must be aligned in the FP register, which are 64-bits wide.

Control Transfer Instructions

Example instruction	Instruction name	Meaning
<code>jal x1,offset</code>	Jump and link	$\text{Regs}[x1] \leftarrow \text{PC}+4; \text{PC} \leftarrow \text{PC} + (\text{offset} \ll 1)$
<code>jalr x1,x2,offset</code>	Jump and link register	$\text{Regs}[x1] \leftarrow \text{PC}+4; \text{PC} \leftarrow \text{Regs}[x2] + \text{offset}$
<code>beq x3,x4,offset</code>	Branch equal zero	$\text{if} (\text{Regs}[x3] == \text{Regs}[x4]) \text{PC} \leftarrow \text{PC} + (\text{offset} \ll 1)$
<code>bgt x3,x4,name</code>	Branch not equal zero	$\text{if} (\text{Regs}[x3] > \text{Regs}[x4]) \text{PC} \leftarrow \text{PC} + (\text{offset} \ll 1)$

Figure A.27 Typical control flow instructions in RISC-V. All control instructions, except jumps to an address in a register, are PC-relative.


RISC-V Dynamic Instruction Mix for SPECint2006

Program	Loads	Stores	Branches	Jumps	ALU operations
astar	28%	6%	18%	2%	46%
bzip	20%	7%	11%	1%	54%
gcc	17%	23%	20%	4%	36%
gobmk	21%	12%	14%	2%	50%
h264ref	33%	14%	5%	2%	45%
hammer	28%	9%	17%	0%	46%
libquantum	16%	6%	29%	0%	48%
mcf	35%	11%	24%	1%	29%
omnetpp	23%	15%	17%	7%	31%
perlbench	25%	14%	15%	7%	39%
sjeng	19%	7%	15%	3%	56%
xalancbmk	30%	8%	27%	3%	31%

Figure A.29 RISC-V dynamic instruction mix for the SPECint2006 programs. Omnetpp includes 7% of the instructions that are floating point loads, stores, operations, or compares; no other program includes even 1% of other instruction types. A change in gcc in SPECint2006, creates an anomaly in behavior. Typical integer programs have load frequencies that are 1/5 to 3x the store frequency. In gcc, the store frequency is actually higher than the load frequency! This arises because a large fraction of the execution time is spent in a loop that clears memory by storing x0 (not where a compiler like gcc would usually spend most of its execution time!). A store instruction that stores a register pair, which some other RISC ISAs have included, would address this issue.

RISC-V Hybrid Instruction Encoding

- 16, 32, 48, 64, ... bits length encoding
- Base instruction set (RV32) always has fixed 32-bit instructions with lowest two bits = 11_2
- All branches and jumps have targets at 16-bit granularity (even in base ISA where all instructions are fixed 32 bits)



Byte Address: base+4 base+2 base

Four Core RISC-V Instruction Formats

- <https://github.com/riscv/riscv-opcodes/>

Additional opcode bits/immediate

Additional opcode bits

7-bit opcode field (but low 2 bits = 11_2)

Reg. Source 2

Reg. Source 1

Destination Reg.

31

25 24

20 19

15 14

12 11

7 6

0

funct7

rs2

rs1

funct3

rd

opcode

R-type

imm[11:0]

rs1

funct3

rd

opcode

I-type

imm[11:5]

rs2

rs1

funct3

imm[4:0]

opcode

S-type

imm[31:12]

rd

opcode

U-type


Aligned on a four-byte boundary in memory. There are variants!

Sign bit of immediates always on bit 31 of instruction. Register fields never move.

RISC-V Encoding Summary

Name	Field						Comments
Field size	7 bits	5 bits	5 bits	3 bits	5 bits	7 bits	
R-type	funct7	rs2	rs1	funct3	rd	opcode	Arithmetic instruction format
I-type	imm[11:0]		rs1	funct3	rd	opcode	Loads & immediate arithmetic
S-type	imm[11:5]	rs2	rs1	funct3	imm[4:0]	opcode	Stores
B-type	imm[12,10:5]	rs2	rs1	funct3	imm[4:1,11]	opcode	Conditional branch format
J-type	imm[20,10:1,11,19:12]				rd	opcode	Unconditional jump format
U-type	imm[31:12]				rd	opcode	Upper immediate format

Immediate Encoding Variants



- S-type vs. B-type

- The 12-bit immediate field is used to encode branch offsets in multiples of 2 in the B format. Instead of shifting all bits in the instruction-encoded immediate left by one in hardware as is conventionally done, the middle bits (imm[10:1]) and sign bit stay in fixed positions, while the lowest bit in S format (inst[7]) encodes a high-order bit in B format.

- U-type vs. J-type

- Similarly, the 20-bit immediate is shifted left by 12 bits to form U immediates and by 1 bit to form J immediates. The location of instruction bits in the U and J format immediates is chosen to maximize overlap with the other formats and with each other.

Immediate Encoding Variants


31	30	25	24	21	20	19	15	14	12	11	8	7	6	0	
funct7					rs2			rs1	funct3		rd		opcode		R-type
imm[11:0]							rs1	funct3		rd		opcode		I-type	
imm[11:5]					rs2			rs1	funct3		imm[4:0]		opcode		S-type
imm[12]	imm[10:5]				rs2			rs1	funct3		imm[4:1]	imm[11]	opcode		B-type
imm[31:12]										rd		opcode		U-type	
imm[20]	imm[10:1]				imm[11]			imm[19:12]		rd		opcode		J-type	





31	30	20	19	12	11	10	5	4	1	0			
— inst[31] —						inst[30:25]		inst[24:21]		inst[20]		I-immediate	
— inst[31] —						inst[30:25]		inst[11:8]		inst[7]		S-immediate	
— inst[31] —						inst[7]	inst[30:25]		inst[11:8]		0	B-immediate	
inst[31]	inst[30:20]			inst[19:12]			— 0 —					U-immediate	
— inst[31] —				inst[19:12]			inst[20]	inst[30:25]		inst[24:21]		0	J-immediate

RISC-V Addressing Modes



1. Immediate addressing





2. Register addressing



3. Base addressing, i.e., displacement addressing



4. PC-relative addressing



ALU Instructions: R-Type

R-type (Register)

- rs1 and rs2 are the source register, rd is the destination
- ADD/SUB
- SLT, SLTU: set less than
- SRL, SLL, SRA: shift logic or arithmetic left or right

31	25 24	20 19	15 14	12 11	7 6	0
funct7	rs2	rs1	funct3	rd	opcode	
7	5	5	3	5	7	
0000000	src2	src1	ADD/SLT/SLTU	dest	OP	
0000000	src2	src1	AND/OR/XOR	dest	OP	
0000000	src2	src1	SLL/SRL	dest	OP	
0100000	src2	src1	SUB/SRA	dest	OP	

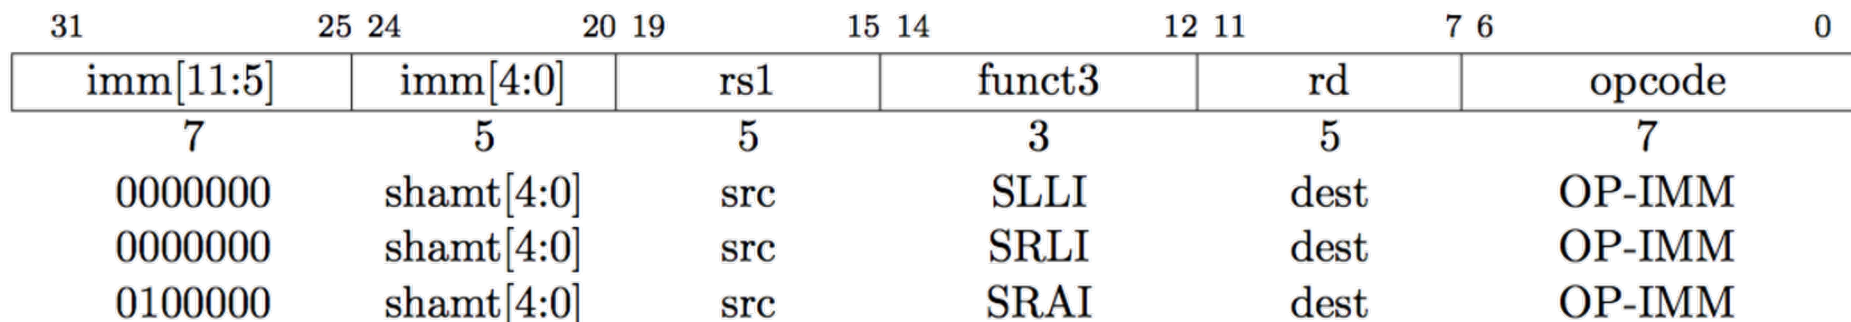
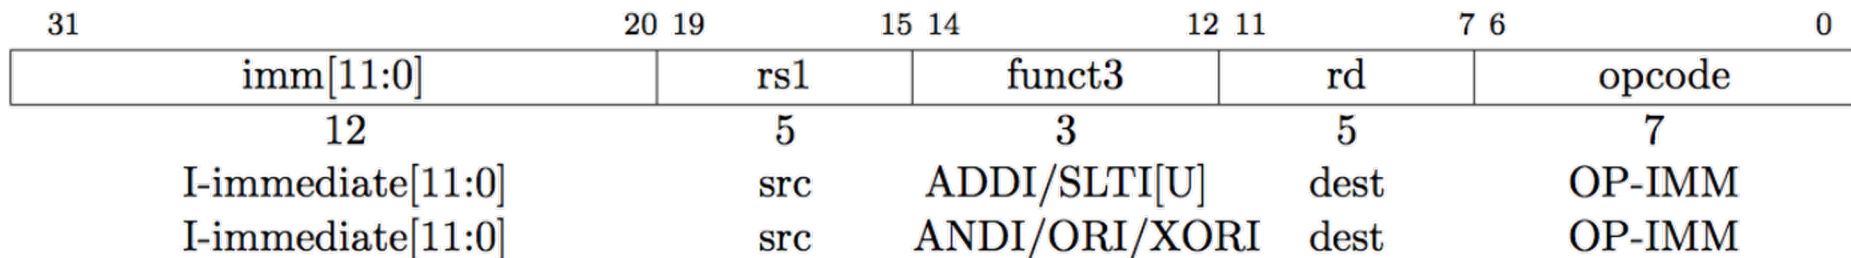
Inst	Name	FMT	Opcode	funct3	funct7	Description (C)	Note
add	ADD	R	0110011	0x0	0x00	rd = rs1 + rs2	
sub	SUB	R	0110011	0x0	0x20	rd = rs1 - rs2	
xor	XOR	R	0110011	0x4	0x00	rd = rs1 ^ rs2	
or	OR	R	0110011	0x6	0x00	rd = rs1 rs2	
and	AND	R	0110011	0x7	0x00	rd = rs1 & rs2	
sll	Shift Left Logical	R	0110011	0x1	0x00	rd = rs1 << rs2	
srl	Shift Right Logical	R	0110011	0x5	0x00	rd = rs1 >> rs2	
sra	Shift Right Arith*	R	0110011	0x5	0x20	rd = rs1 >> rs2	msb-extends
slt	Set Less Than	R	0110011	0x2	0x00	rd = (rs1 < rs2)?1:0	
sltu	Set Less Than (U)	R	0110011	0x3	0x00	rd = (rs1 < rs2)?1:0	zero-extends

ALU Instructions: I-Type

- I-type (immediate), all immediates in all instructions are sign extended

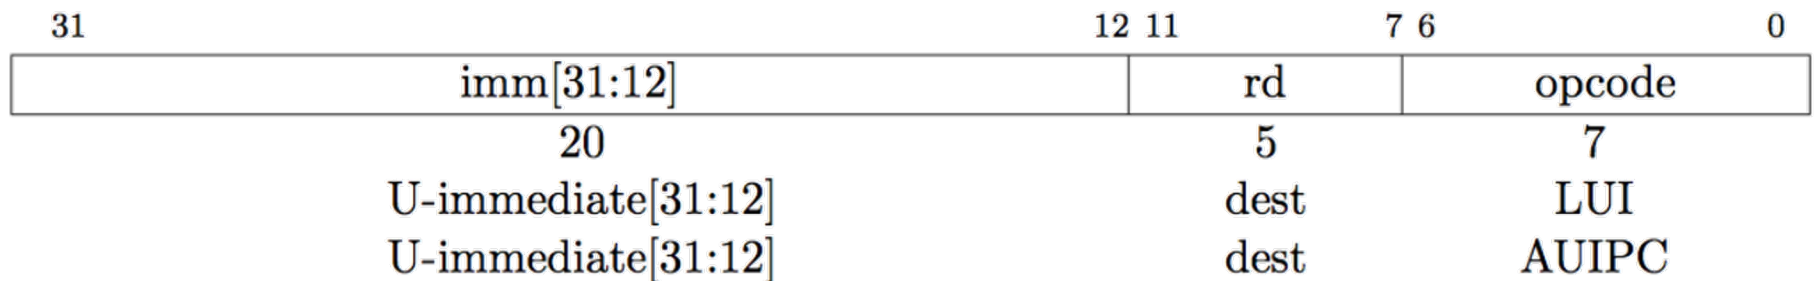
- ADDI: adds sign extended 12-bit immediate to rs1
- SLTI(U): set less than immediate
- ANDI/ORI/XORI: logical operations
- SLLI/SRLI/SRAI: shifts by constants

I-type instructions end with I



ALU Instructions: U-Type



- LUI/AUIPC: load upper immediate/add upper immediate to PC



- Writes 20-bit immediate to top of destination register
- Used to build large immediates
- 12-bit immediates are signed, so have to account for sign when building 32-bit immediates in 2-instruction sequence (LUI high-20 bits, ADDI low-12 bits)



Load/Store Instructions: I/S-Type

- Load instruction (I-type)
 - $rd = MEM(rs1 + imm)$
- Store instruction (S-type)
 - $MEM(rs1 + imm) = rs2$



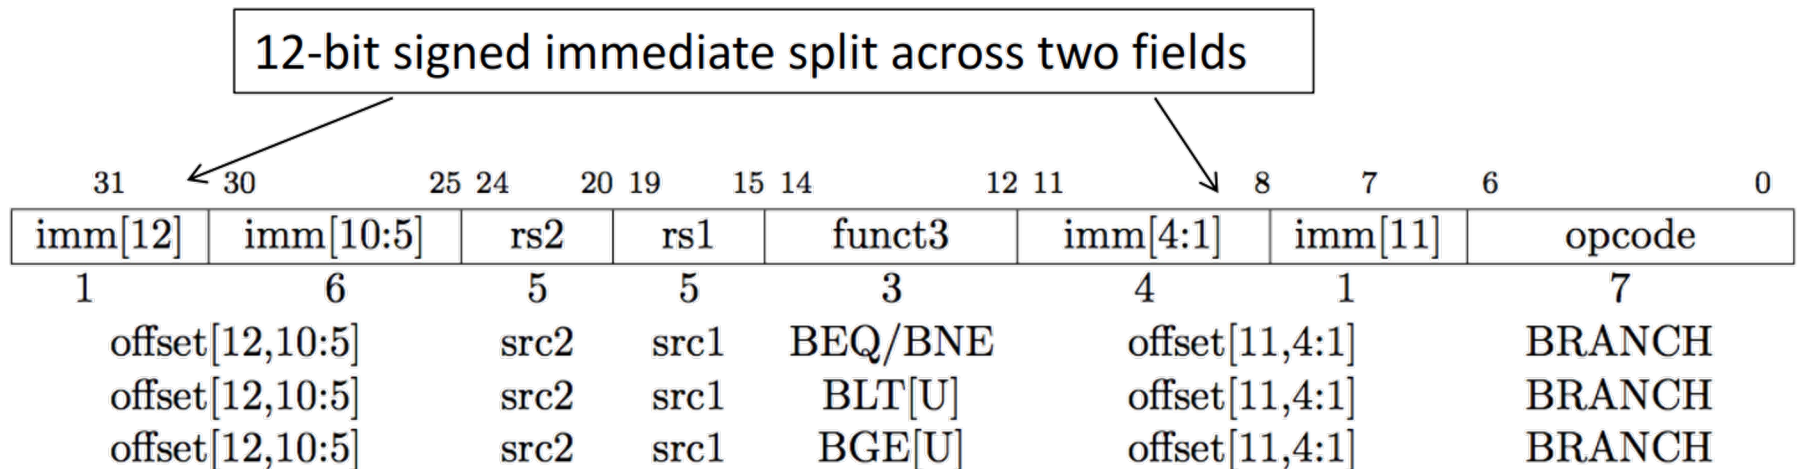
Control Transfer Instructions: J-Type

- No architecturally visible delay slots
- Unconditional jumps: PC + offset target
 - JAL: jump and link, also writes PC + 4 to x1, J-type
 - Offset scaled by 1-bit left shift – can jump to 16-bit instruction boundary (same for branches)
 - JLAR: jump and link register where $\text{imm} (12\text{bits}) + \text{rd1} = \text{target}$



Control Transfer Instructions: B-Type


- No architecturally visible delay slots
- Conditional branches: B-type and PC + offset target



Branches, compare two registers, PC + (immediate << 1) target (signed offset in multiples of two). Branches do not have delay slot.

Where is NOP?

ADDI x0, x0, 0



Privileged ISA: Modes

- RISC-V privileged spec defines 3 levels of privilege, called modes
 - Machine mode is the highest privileged mode and the only required mode

Level	Encoding	Name	Abbreviation
0	00	User/Application	U
1	01	Supervisor	S
2	10	<i>Reserved</i>	
3	11	Machine	M

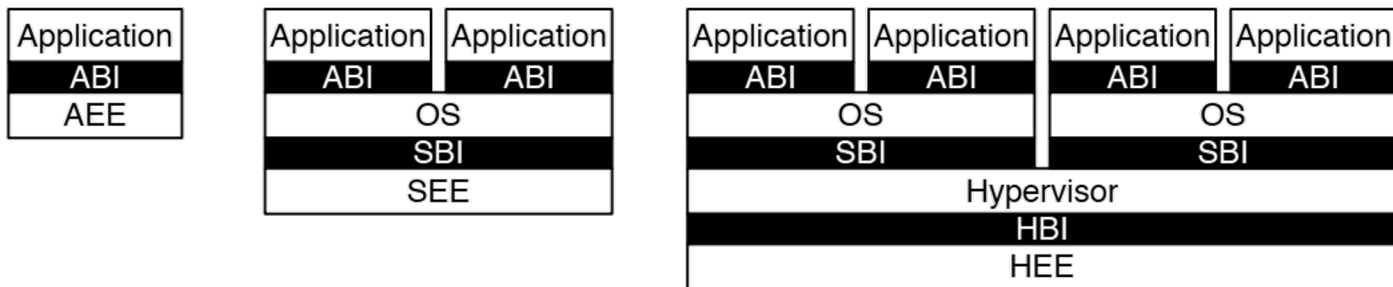
- More-privileged modes generally have access to all of the features of less-privileged modes, and they add additional functionality not available to less-privileged modes, such as the ability to handle interrupts and perform I/O. Processors typically spend most of their execution time in their least-privileged mode; interrupts and exceptions transfer control to more-privileged modes.

Software Stack and Instructions

- Implementations might provide anywhere from 1 to 4 privilege modes trading off reduced isolation for lower implementation cost

Number of levels	Supported Modes	Intended Usage
1	M	Simple embedded systems
2	M, U	Secure embedded systems
3	M, S, U	Systems running Unix-like operating systems

- RISC-V privileged software stack



- RV32/64 privileged instructions

$\left. \begin{array}{l} \underline{\text{machine-mode}} \\ \underline{\text{supervisor-mode}} \end{array} \right\} \text{trap } \mathbf{return}$
 $\underline{\text{supervisor-mode}} \mathbf{fence} \mathbf{virtual\ memory\ address}$
 $\underline{\text{wait}} \mathbf{for} \mathbf{i} \underline{\text{interrupt}}$

Specifications and Software

- Specification from RISC-V website
 - <https://riscv.org/specifications/>
- RISC-V software includes
 - Toolchain projects
 - <https://wiki.riscv.org/display/HOME/Toolchain+Projects>
 - A simulator (“spike”)
 - <https://github.com/riscv-software-src/riscv-isa-sim>
 - Standard simulator QEMU (Upstream now)
 - <https://github.com/riscv/riscv-qemu>
- Operating systems support exists for Linux (Upstream now)
 - <https://github.com/riscv/riscv-linux>
- A javascript ISA simulator to run a RISC-V Linux system on a web browser
 - <https://github.com/riscv/riscv-angel>