

In this lecture, the instruction set architecture (ISA) of the RISC-V processor will be introduced. We will only consider the base instruction set for the 32-bit integer version of the ISA. Focus will be on the six different instruction types with emphasis each instruction's functionality and encoding of its machine code.

ISA of a processor does not dictate how the processor is implemented. It only defines how to the processor is programmed. The actual hardware architecture will be covered in the next lecture.



RISC-V is by no means the first or the only Reduced Instruction Set Computer that enjoys widespread adoption. The early RISC processor from Berkeley, the MIPS, was used for over four decades in industry. The UK's ARM processor (original stands for Acorn Risc Machine) is the most manufactured CPU in history with an estimated 200 billions being shipped to date. However, RISC-V is the first widely accepted opensource RISC processor. Opening the ISA to the public (royalty free) is a new business model and has captured much attention in the past 5 years. Together with the free toolchain for development and many open-source design freely available, RISC-V is expected to pose real competition to x86 and ARM architecture to become at least one of the dominant play in the sector.

RISC-V was developed by Krste Asanovic and Dave Patterson (and others) in Berkeley in early 2010's. The ISA was first published in 2011, and its future development and ratification are under the control of RISC-V Foudation and RISC-V International located in Switzerland.

According to Patterson and Hennessy's textbook, the underpinning design principles are shown on the slides.

# **Design Principles**

### Principle 1: Simplicity favors regularity

- Consistent instruction format
- Same number of operands (two sources and one destination)
- Easier to encode and handle in hardware

#### Principle 2: Make the common case fast

- RISC-V includes only simple, commonly used instructions
- · Hardware to decode and execute instructions can be simple, small, and fast
- More complex instructions (that are less common) performed using multiple simple instructions
- RISC-V is a reduced instruction set computer (RISC), with a small number of simple instructions
- Other architectures, such as Intel's x86, are complex instruction set computers (CISC)

## **Principle 3: Smaller is Faster**

H&H p301-303

PYKC 29 Oct 2024

EIE2 Instruction Architectures & Compilers

Lecture 6 Slide 3

As shown in later slides, the instruction set has highly consistent format. The base instruction set is known as RV32I (RISC-V 32-bit integer only) only has 40 instructions. The ISA has two sources and one destination operands. The format of the instructions are divided into only six different types.

The second principle is the RISC-V only has commonly used instructions. Yet, it is Turing Complete, meaning that it can be used to implement any computer algorithms. This makes the RISC-V implementation both small and fast. Complex operations are achieved by stringing together multiple instructions.



To understand the RV32I ISA, we start with the add and subtract instructions as shown in this slide. The diea is very simple – it shows that for add instructions, we need two source operands (b and c) and one destimation (a).



Where do operands come from and where does the destination operand go?

There are three possibilities. The fastest and most often used operand is from or to **registers** on the CPU chip itself. RISC-V RV32I has 32 32-bit registers. They form an integral part of the CPU design and accessing them is easy and fast. 32 registers means the instruction must use  $3 \times 5$ -bit = 15 bit of the 32-bit instruction to specify a register-only (R-type) instruction.

The second possibility is from data memory. To access this, the instruction must specify the data memory address, using a register as a pointer. In RISC-V, the register content is often specified with an associated offset constant value as part of the instruction.

The third possibility is from instruction memory, i.e. the operand is a constant within the instruction itself. In "RISC-V speak", this is called an immediate.

Instruction Type	31	30 2!	J 28	27 2	5 25	24	23	22 2	1 2	0	19 18	17	16 15	14	13 1	.2 1	10	9	8	7	6	5	4	3	2	
Register/register			funct	t7				rs2		t		rs1		fu	nct3			rd					ор	cod	le	
Immediate (I-type)	-			imn	[11:	0]	~			T		rs1		fu	nct3			rd					ор	cod	le	
Upper (U-type)	e (U-type) imm[11:5]								:12]									rd					ор	cod	e	
Store (S-type)	e (S-type) imm[11:5] th (B-type) [12] imm[10:5]							rs2				rs1		fu	nct3		im	m[4	:0]				ор	cod	le	
Branch (B-type)								rs2				rs1		fu	nct3		imm	[4:1	]	[11]			ор	cod	le	
Jump (J-type)	[20]			imn	n[10:	:1]			[1	1]		im	m[19:	12]				rd					ор	coc	le	
	o (J-type) [20] imm[10: code (7 bit): partially specifi nct7 + funct3 (10 bit): comb 1 (5 bit): specifies register co																									
<ul> <li>opcode (7</li> <li>funct7 + fu</li> <li>rs1 (5 bit):</li> <li>rs2 (5 bit):</li> </ul>	nct spe	: par 3 (10 cifie	tiall ) bi s re s se	t): co	ecif omb er co d re	ies oine ont	ed v taini	vith ing t	of t op first	t o	e 6 ty ode, opera	/pe: the	s of <i>i</i> se tw	nstr o fi	elds	on s de	forn	ibe	s wł	nat	op	era	atio	on t	o pe	erfo

RISC-V RV32I has six types of instructions.

**R-type** (**Register/register**) instructions use only registers as source and destiantions. This instruction type is mostly used for arithmetic and logic operations involving the ALU.

**I-type (Immediate)** instructions has one of the two source operands specified within the 32-bit instruction word as a 12-bit constant (or immediate). This constant is regards as 12-bit signed 2's complement number, which is always sign extended to form a 32-bit operand.

**S-type (Store)** instructions are exclusively used for storing contents of a register to data memory.

**B-type (Branch)** instructions are used to control program flow. It compares two operands stored in registers and branch to a destination address relative to the current Program Counter value.

J-type (Jump) instructions are used for subroutine calls.

**U-type (Upper immediate)** instructions are used to specify the upper 20 bits immediate value of a register.

Name	Register Number	Usage
zero	x0	Constant value 0
ra	x1	Return address
sp	x2	Stack pointer
gp	x3	Global pointer
tp	x4	Thread pointer
t0-2	x5-7	Temporaries
s0/fp	x8	Saved register / Frame pointer
s1	x9	Saved register
a0-1	x10-11	Function arguments / return values
a2-7	x12-17	Function arguments
s2-11	x18-27	Saved registers
t3-6	x28-31	Temporaries

RISC-V RV32 has 32 registers designated as **x0** to **x31**. They are "general purpose" registers in the sense that the ISA allows them to be used for any purpose with the exception of **x0**, which ALWAYS contain the value 32'b0. Writing to x0 does not change its content.

Having **x1** to **x31** for any general use can be confusing. Common good practice is included in a guideline where specific registers are used for special functions. For example x1 is used to store the **return address** (of a subroutine) and therefore **x1** is also called **ra**.

The table above shows the various aliases for all 32 registers. You are recommend o use the given name of these registers to make the program more readable. For example instead of using **x0**, you should always refer to it as **zero**.

		N	am	e	R	legi	ste	er N	umb	ber	U	sag	e										
		s	)/f	p	X	8					S	aveo	d re	giste	r / F	ram	e p	oint	er				
		s	1		x	9					S	aveo	d re	giste	r								
		s	2-1	.1	x	18-	27				S	aveo	d re	giste	rs								
C	Co	ode								RI	SC	-V a	iss	emb	oly	cod	е					-	
										#	s0	= a	a,	s1 =	= b	, s	2 =	= C					
a	=	b +	с;							ad	ld	s0,	s1	, sź	2								
a	=	b +	6;							#	s0	= ;	a,	s1 =	= b								
51 044										au	iur	50,	,	•±,	5								_
struction ormats	31	30 29	28	27 26	25	24	23 2	22 21	20	19	18	17 16	15	14 13	12	11 1	9	8 7	6	6 5	4 3	3 2	1
egister/register		f	unct7				r	s2				rs1		func	t3		rd				орс	ode	
mediate				imm[	11:0]						)	rs1		func	t3		rd				орс	ode	

Consider again the add instructions. ADD is a typical ALU instruction in the class of arithematic and logic operations. It needs two source operands and one destination operands to store the results. Shown here is the instruction: add s0, s1, s2 which uses three registers. Consider the encode of this instructions (slide 21).

ор	funct3	funct7	Туре	Instruction	Description	Operation
0110011 (51)	000	0000000	R	add rd, rs1, rs2	add	rd = rs1 + rs2

The operation is specified with the opcode, funct3 and funct7 fields of the instructions. opcode = 7'h38 (51), funct3 = 3'b0, funct7 = 7'b0.

rd: s0 = x8 = 5'b01000, rs1 = s1 = x9 = 5'b01001, rs2 = s2 = x18 = 5'b10010If we fill in the fields with these values according the diagram here, we get:

Instruction Formats	31	30	29	28	27	26	25	24	23	22	21	20	19	18	17	16	15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
Register/register			fu	Inct7	7					rs2					rs1			f	unct	3			rd					op	cod	le		
		000	00		0	00		1		00	10			01	00		1		00	0		010	00		0	0	11		0	00	11	

Therefore this instruction has a machine code of 32'h01248433.

Similar, for: addi s0, s1, 6

ор	funct3	funct7	Туре	Instruction	Description	Operation
0010011 (19)	000	-	Ι	addi rd, rs1, imm	add immediate	rd = rs1 + SignExt(imm)

opcode = 7'h13 (19), funct3 = 3'b0.

rd: s0 = x8 = 5'b01000, rs1 = s1 = x9 = 5'b01001 as before.

 $Imm_{12} = 12'h6.$ 

Instruction Formats	31	30	29	28	27	26	25	24	23	22	21	20	19	18	17	16	15	14	13	12	11	10	9	8	7	6	5	4	3	2 1	1	)
Immediate						imm	[11:0	]							rs1			1	funct	3			rd					opo	cod	Э		
		000	00			00	00			01	10			01	00		1		00	0		010	00		0	0	01		0	01	1	

Therefore this instruction has a machine code of 32'h00648413.



32-bit operands in memory occupies 4 bytes. Some processor uses one unique address for each 32-bit words. MIPS processor is one such example. Everything instruction or data word has a unique address – it is "word addressable" processor.

RISC-V uses byte-addressable to access memory, where EVERY BYTE has a unique address.



Therefore, in RISC-V, every 32-bit occupies four unique addresses. If the least significant byte has an address of base = 4, then the most significant byte has an address of base + 3 = 7 as shown above.

Since all RISC-V instructions are 32-bit, addresses of the instruction memory are all aligned to an increment of 4.



This shows an example of how reading from data memory into register s3 at data memory address 32'h8.



This is an example of storing from register t7 to memory address 32'h10.



An operand can also be a constant encoded within the instruction itself. Here comes a problem: since RISC-V instructions are all **single** 32-bit words, and an operand is also 32-bit wide, how can an immediate constant operand be embedded in a 32-bit instruction?

If the constant operand has a value of -2048 to -2047 (12'hFFF to 12'h8FF), the operand can be fully specified with a 12-bit binary number in 2's complement form. As it turns out, most constants in computer programs are small. For example, to refer to an offset index of an array, the index often falls within this range of numbers.

In RV32I, I-type instructions have 12 bits reserved for such a constant operand as shown in the slide here. The constant is always **sign externded** before being used as an operand.

		F	RI	S	C-	V		0	)p	e	ra	n	d	V	vi	th	3	<b>32</b>	-k	bit	t (	Co	or	າຮ	ta	ar	nt	S					
•	Us lui an	e և ։ բւ d 0	oad uts 's i	d u ai n	ıpı n i lov	oei mi we	r ir me r 1	mn edi L2	ne ate bit	dia e i ts	ate n t	e ( :h	(lu ie i	ıi) up	ar op	nd er	ас 2С	ldi ) b	its	5 C	ofo	de	st	ina	iti	or	ר ר	eg	gist	er			
	СС	Cod	le														R #	RIS	<b>C</b> .	-V	as	S	ən	ıbl	у	cc	bd	е					
	int	t a	=	0:	xFl	EDO	281	765	;								" l a	ui dd	i	s( s(	), ),	0: s	xF O,	ED( 0:	C8 x7	65	5						
	Re	me	m	be	rt	ha	t a	ado	di s	sig	;n-	e	xt	en	d	s it	S 2	12 <sup>.</sup>	-b	it	im	m	e	dia	te	e c	01	nst	tar	t			
Instruction Formats		31	30	29	28	27	26	25	24	23	22	2	1 2	20	19	18	17	16	15	5 1	4 1	3	12	11	10	9	8	7	6	5 4	3	2 1	0
Upper Immediate										i	mm[	31:	:12]													rd				o	bcod	е	
PYKC 2	29 Oct 2	2024									EIE2	2 In:	struc	ctior	Arc	chited	cture	s & (	Corr	npile	rs								Ha	&H p ure 6	306 Slide	ə 14	

Using a 12-bit immediate constant works most of the time. However, there are times when a program requires to load a register (say) with a 32-bit constant value.

In RV32I, this is achieved by splitting the constants into two parts – the upper 20 bit, which can be loaded into a register, using the instruction "load upper immediate" lui.

For example, the instruction:  $lui \pm 0$ ,  $0 \times FEDC8$  load into s0 the value 32'hFEDC8000. This is then added to the bottom 12 bits of the constant with the "add immediate" addi instruction:

addi s0, s0 0x765.

This works perfectly if the MSB of the 12-bit immediate operand is 0. Unfortunately, if the MSB of the 12-bit constant (i.e. bit 11) is a 1, the constant is then sign extended. When added to the upper 20-bits previously loaded value in s0, the answer will be wrong because the upper 20-bit will be modified. This is because in 2's complement representation, a 20-bit value of 20'hFFFFF is equivalent to -1. Therefore the upper 20-bit, after the addi instruction with be 1 lower than what it should be.

bit 11 of the constar	nt is <mark>1</mark> , increment upper 20 bits by <mark>1</mark> in lui
<b>C Code</b> int a = 0xFEDC <mark>8E</mark> AB;	<b>Note:</b> -341 = 0xEAB
RISC-V assembly cod	le
lui s0, 0xFEDC <mark>9</mark>	# s0 = 0xFEDC9000
addi s0, s0, -341	<pre># s0 = 0xFEDC9000 + 0xFFFFFEAB</pre>
	# = 0xfedc8eab

Therefore, if bit 11 of the 32-bit constant is 1, we load the upper 20-bit with a constant that is 1 larger than the constant.

In this example, the constant is 32'hFEDC8EAB. Bit 11 is 1. Upper 20-bit is 20'hFEDC8, and lower 12-bit is 12'hEAB, which is -341 in 2's complement representation after sign extension.

We first load s0 with 0xFEDC9 (1 larger than the upper value). After the addi instruction, s0 will have the correct 32-bit constant value.

Fortunately the assembly and compiler for RISC-V take care of this automatically.



RISC-V has many instructions missing deliberately to make is small and fast. More complex operations are accomplished by multiple instructions or by an instruction that result in the same operation.

For example there is no instruction to load a register with a constant value. To load s0 with the small constant 6, we use the instruction:

addi s0, zero, 6

To load s0 with a large constant 0xFEDC8EAB, we use the two instructions:

lui s0, 0xFEDC9
addi s0, s0, 0xEAB

This makes the assembly language program of RISC-V much harder to read and understand. Fortunately, RISC-V assembler understand a number of pseudo instructions. These instructions do not exist in the RISC-V ISA, but are translated into equivalent RV32I instructions.

To load a register with a constant of any size constant (up to 32 bits), one can use the "load immediate" li pseudoinstruction.

li s0, 6 li s0, 0xFEDC9

Slide 29 shows all the pseudo instructions that RISC-V assembler accepts.



Specifying where the operand comes from is called "addressing modes" of an ISA. We have already discussed the two of the four addressing modes found in RISC-V ISA: Register addressing and Immediate addressing. We will now consider the remain two other addressing modes: Base addressing (with offset) and Program Counter Relative addressing.



**Base addressing** mode uses one of the registers content as the address into memory. What stored in the register is not the actual operand, but it stores the address of the operand. In C++, we call this a **pointer** - it points to the place where the operand is stored.

In RISC-V, Base addressing is always used with an offset value which must be a 12-bit 2's complement immediate constant. The "load" and "store" instrutions use this mode of addressing.

Example:	_	
<b>Address</b> 0x354 0x358	L1:	Instruction addi s1, s1, 1 sub t0, t1, s7
 OxEBO		 bne s8, s9, L1
The label is (0x	EBO-0x35	54) = 0xB5C ( <b>2908</b> ) instructions <b>before</b> br
		54) – 0xb3C ( <b>2508</b> ) instructions <b>before</b> bi

The final addressing mode is the Program Counter, or PC-relative addressing. The operand is derived from the PC value by adding a 13-bit (not 12-bit) 2's complement offset. This type of addressing is ONLY used by the branch and jump instructions.

For example, the above "branch if not equal" instruction compares s8 and s9 contents. If they are NOT the same, then the PC counter is load with the address of L1, which is 0x354.

How is the value 0x354 encoded in the instruction? The immediate constant is calculated with the value of PC for the bne instruction, which ix 0xEB0. The offset is calculated by 0xEB0 - 0x354 = 0xB5C. Therefore the stored immediate value is therefore the value -2908.



The way that RISC-V encodes the relative offset of -2908 is complicated and appears illogical. In fact the design decision for this instruction is very clever and is aimed at making the hardware implementation as simple as possible. Here are the design constraints that determine how the instruction is encoded:

- 1. It uses the same fields for opcode (7 bits), funct3 (3 bits), rs1 and rs2 (5 bits) as other instructions. This means that 20 bits of the 32-bit instructions are already used. So there are 12 bits left for encoding the offset.
- 2. Since the branch destination is ALWAYS an instruction address, and that RISC-V uses byte-addressable memory, the instructions for RV32I is ALWAYS aligned to 4. In other words, there is no need to store the bottom 2 bits of the offset they are always zero. However, there is a variant of RISC-V ISA which targets microcontroller, where the among of program memory is limited. The "Compressed" extension of RISC-V ISA includes 16-bit instructions (i.e. packing two instructions into a 32-bit word). Therefore, the instruction address can be an increment of 2 instead of 4, meaning that only bit 0 is always 0.
- 3. It is convenient in hardware that the bits used for encoding B-type immediate values should be similar to that used for I-type and S-type instructions. Therefore the locations of bits are the same for imm[4:1], imm[10:5]. However, the branch immediate is 13 bits instead of 12 bits, therefore imm[12] now takes the place of imm[11] in other case. They are both sign bits.
- 4. Since imm[0] is always 0, there is no need to store it. Instead imm[11] is

stored here!

Instruction Formats	31	30 29	28	27 2	5 25	2	24 23	22	21 20	19	18	17	16	15	14	13	12	11	10	9 8	7	6	5 5	4	3	2 1	1
Register/registe	r	f	unct7	,				rs2				rs1			fu	Incta	3			rd				0	pcod	е	
op	funct3	funct7		Туре	Instr	uct	ion			Desc	ripti	on					Op	рега	tion								
0110011 (51)	000	00000	00	R	add	ľ	rd, ı	rs1,	rs2	add							rd	=	rs1	+	rs	52					
0110011 (51)	000	01000	00	R	sub	ľ	rd, 1	rs1,	rs2	sub							rd	-	rs1	-	rs	52					
0110011 (51)	001	00000	00	R	s11	ľ	rd, ı	rs1,	rs2	shift	left	logic	al				rd	=	rs1	<<	rs	24:0					
0110011 (51)	010	00000	00	R	slt	1	rd, ı	rs1,	rs2	set le	ess th	an					rd	=	(rsl	<	rs	2)					
0110011 (51)	011	00000	00	R	sltu	1	rd, 1	rs1,	rs2	set le	ess th	an u	insig	ned	2		rd	-	(rs]	<	rs	2)					
0110011 (51)	100	00000	00	R	xor	1	rd, ı	rs1,	rs2	xor							rd	=	rs1	^	rs	2					
0110011 (51)	101	00000	00	R	srl	ľ	rd, ı	rs1,	rs2	shift	righ	t log	ical				rd	=	rs1	$\rightarrow$	rs	24:0	6				
0110011 (51)	101	01000	00	R	sra	1	rd, ı	rs1,	rs2	shift	righ	t ari	thme	tic			rd	-	rs]	>>	> rs	24:0	į				
0110011 (51)	110	00000	00	R	or	ľ	rd, 1	rs1,	rs2	or							rd	=	rs1		rs	2					
0110011 (51)	111	000000	00	R	and	1	rd, ı	rs1,	rs2	and							rd	=	rs1	&	rs	2					

This an the next few slides are summary of ALL the 40 instructions in RISC-V RV32I ISA.

Here is the R-type instructions that perform arithmetic and logical operations using three registers. They all share the opcode of 51 decimal (or 0x33). The funct3 and funct7 fields defines the specific operation.

Instruction Formats	31	30	29	28	27	26	25	24	23	22	21	20	19	18	7 1	16	15	14	13	12	11	10	9	8	7	6	5	4 3	3 2	1
Immediate					i	mm	[11:0]							r	51			fi	unct	3			rd					орс	ode	
Store			im	n[11	:5]					rs2	2			r	s <b>1</b>			f	unct	3		im	m[4:	0]				opc	ode	
ор	funct3	l fu	nct7		Typ	e I	nstru	ictio	n				Desc	riptio	n					0	рега	tion								
0000011 (3)	000	-			I	1	b	rc	١,	imm	(rs1	.)	load	byte						rc	i =	Si	gnEx	t(	[Add	dre	ss];	7:0)		
0000011 (3)	001	-			I	1	h	rc	١,	imm	(rs1	.)	load	half						ro	i =	Si	gn E x	t(	[Add	dre	ss]	15:0	)	
0000011 (3)	010	-			I	1	W	rc	١,	imm	(rs1	)	load	word						rc	i =				[Add	dre	ss]	31:0		
0000011 (3)	100	-			I	1	bu	rc	١,	imm	(rs1	)	load	byte	insi	gnee	d			ro	1 =	Ze	°0Ex	t(	[Add	dre	ss]	7:0)		
0000011 (3)	101	-			Ι	1	hu	rc	١,	imm	(rsl	)	load	half u	nsig	gnec	1			rc	i =	Ze	^oEx	t(	[Add	ire	ss]	15:0	)	
0010011 (19)	000	-			I	ő	ddi	rc	١,	rs1	, im	ım	add	imme	liate	e				ro	=	rs	+		Sign	nEx	t(in	mm )		
0010011 (19)	001	00	000	$00^{*}$	Ι	5	:11i	rc	١,	rs1	, ui	mm	shift	left le	gica	ıl in	nme	diat	e	rc	1 =	rs	<<		uimr	n				
0010011 (19)	010	-			I	5	ilti	rc	١,	rs1	, im	ım	set le	ess tha	n in	nme	edia	te		rc	1 =	(rs	<		Sign	۱Ex	t(ir	nm )	)	
0010011 (19)	011	-			Ι	5	iltiu	u ro	Ι,	rs1	, im	ım	set le	ess tha	n in	nm.	uns	signe	ed	rc	i =	(rs	<		Sign	۱Ex	t(ir	nm )	)	
0010011 (19)	100	-			Ι	>	ori	rc	١,	rs1	, im	ım	xor	imme	liate					ro	i =	rs	^		Sig	ηEx	t(in	mm )		
0010011 (19)	101	00	000	$00^{*}$	Ι	5	rli	rc	,	rs1	, ui	mm	shift	right	logi	cal i	imn	nedia	ate	rc	i =	rs	L >>	,	uimr	n				
0010011 (19)	101	01	000	$00^{*}$	Ι	5	srai	rc	١,	rs1	, ui	mm	shift	right	aritl	hme	etic	imm	۱.	rc	i =	rs	1 >>	>	uimr	n				
0010011 (19)	110	-			Ι	0	ori	rc	١,	rs1	, im	ım	or in	nmedi	ate					rc	1 =	rs	L		Sig	nEx	t(in	mm )		
0010011 (19)	111	-			Ι	õ	ndi	rc	,	rs1	, im	ım	and	imme	liate	•				rc	i =	rs	L &		Sig	ηEx	t(in	mm )		
0100011 (35)	000	-			S	5	sb	rs	2,	imm	(rs1	)	store	e byte						[ A	١ddr	ess	]7:0	=	rs2;	7:0				
0100011 (35)	001	-			S	5	sh	rs	2,	imm	(rs1	)	store	e half						[A	∖ddr	ess	]15:0	-	rs2	5:0				
0100011 (35)	010	-			S		ŚW	rs	2,	imm	(rsl	)	store	e wor						[A]	١ddr	ess	]31:0	=	rs2					

This group includes two instruction types which both require TWO register operands and one 12-bit immediate operands.

Instruction Formats	31	30	29	28	27	26	25	24	23	22	21	20	19	18	17	16	15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	D
Immediate					i	imm[	11:0	]							rs1			f	unct	3			rd					ор	cod	е		
Store			imr	n[11	:5]					rs2					rs1			1	funct	3		im	m[4	:0]				ор	cod	е		

The I-type instructions specify either a load instruction or a ALU instructions. Here we specify a destination register rd to store the result of an memory read of the ALU operation, and a source register rs1 to specify an operand for the ALU operation or the address of the data to fetch.

Two opcodes are used for I-type instructions: 3 for load instructions and 19 for ALU immediate instruction. Note that some I-type instructions (shift instructions) do not use sign-extension to the immediate values.

The S-type instructions does not require a destination register because the destination is data memory. However they require two source registers, one contain the value to write to memory, and a second has the base address of the destination. The 12-bit immediate offset is split into two parts, using the funct7 field of instr[31:25] and the rd field of instr[11:7], combined to form imm[11:0].

	B-	type	Ins	str	ucti	0	ns:	Ρ	PC	-re	ela	ati	iv	e	B	sr	ar	າດ	:ł	1e	S					
Instruction Formats	31	30 29 28	27 26	6 25	24 23	22	21 20	19	18	17 1	16	15	14	13	12	11	10	9	8	7	6	5	4	3 2	2 1	0
branch	funct2	funct7	Tupe	Inste	uction	152		Dec	crint	ion			Iu	note		200	rati			[11]			Opt	Joue	,	
op	000	Tunct/	в	hea	rs1	rs2	lahel	bra	nch i	f					i	f	(rs]		= r	c2)	PC	-	RTA			
1100011 (99)	000	-	B	bne	rs1,	rs2,	label	bra	nch i	i = f ≠					i	f	(rs]	≠	r	s2)	PC	=	BTA			
1100011 (99)	100	-	B	blt	rs1.	rs2	label	bra	nch i	f <					i	f	(rs]	<	r	s2)	PC	-	BTA			
1100011 (99)	101	-	B	bge	rs1,	rs2,	label	bra	nch i	f≥					i	f	(rs]	2	r	s2)	PC	=	BTA			
1100011 (99)	110	-	В	bltu	rs1,	rs2,	label	bra	nch i	f < ur	nsig	ned			i	f	(rsl	<	r	s2)	PC	=	BTA			
1100011 (99)	111	-	В	bgeu	rs1,	rs2,	label	bra	nch i	f≥ur	isig	ned			i	f	(rs]	2	r	s2)	PC	=	BTA			
																						Hð	≩H p	311		ĺ
PYKC 29 Oc	t 2024					EIE2	Instructio	on Ar	chited	tures	& C	omp	ilers								L	.ectı	ure 6	i Slie	de 2	3

We have discussed the encoding of branch instructions in details in slide 19 & 20. Note that the opcode for B-type instructions is 99 or 0x63. funct3 defines the conditions under which branch takes place.

When implementing B-type instruction in hardware, one could use the ALU to perform the comparison, or create special branch unit which provides performs ONLY the comparison and no other operations and generates all the required conditions. The second option makes the design cleaner.

	8 U	<u>k</u>	-1	ty	р	e	n	st	ru		tion	S:		U	p	)e	r	&	Jı	u	m	p	/L	.in	ık	(				
Instruction Formats	31	30	29	28	27	26	25	24	23	22	21 20	19	18	17	16	15	14	13	12	1	1 10	9	8	7	(	6 5	4	3	2	1 0
Upper Immediate									ir	nm[3	31:12]											ro	i				0	рсос	de	
Jump	[20]				I	imm[	10:1	I			[11]			ir	nm[1	19:12	2]					ro	l		ſ		0	pcoc	le	
						1-								_																
op	tunct3	fur	ict7		Тур	be L	nstru	ictio	n	unim	195	Desc	ript	ion		1		DC	0	pe	ratio	n á m		1.016	0.1		20			
0110111 (23)		-			U	1	uipu	rd nd	, (	up in	um	add	upp	er in	imeo	diate	e to	PC	r	1 = d =	= {up	im	m,	12'0	U}	+ +	~			
1100111 (103)	-	-			T	i	alr	rd	, (	rs1.	imm	ium	upp	d lin	hine k re	riste	e r		P	c =	= (u)	+ 5	in,	IE NEXT	(i	, mm)	. r	d =	PC	+ 4
1101111 (111)	-	-			J	j	al	rd	, .	l a be	1	jumj	o an	d lin	k It.	giste	-1		P	C =	JTA	,					r	-d =	PC	+ 4
• V PYKC 29 Oct	<b>Ve v</b>	vill	d	isc	us	s a	iui	pc,	, ja		and j	al	ins	str	UC	tic	DN	s ir	n a	n	oth	ie	rl	ec	tu	Jre	<b>e</b> 6	Slid	e 24	

Finally there are four special instructions that are not in the other category. We have already discussed the lui instruction previously.

The U-type instructions are used to manipulate the upper 20-bit of a register to handle 32-bit immediate constants.

The J-type instructions are for function or subroutine calls. They will be discussed in a later lecture.

Mnemonic	Instruction	Туре	Description
ADD rd, rs1, rs2	Add	R	rd ← rs1 + rs2
SUB rd, rs1, rs2	Subtract	R	rd ← rs1 - rs2
ADDI rd, rs1, imm12	Add immediate	I.	rd ← rs1 + imm12
SLT rd, rs1, rs2	Set less than	R	rd ← rs1 < rs2 ? 1 : 0
SLTI rd, rs1, imm12	Set less than immediate	Т	rd ← rs1 < imm12 ? 1 : 0
SLTU rd, rs1, rs2	Set less than unsigned	R	rd ← rs1 < rs2 ? 1 : 0
SLTIU rd, rs1, imm12	Set less than immediate unsigned	I.	rd ← rs1 < imm12 ? 1 : 0
LUI rd, imm20	Load upper immediate	U	rd ← imm20 << 12
AUIP rd, imm20	Add upper immediate to PC	U	rd ← PC + imm20 << 12

EIE2 Instruction Architectures & Compilers

Lecture 6 Slide 25

The next few slides provide a catalogue of all the RISC-V RV32I instructions in various groups.

All these instructions involve arithmetic operation.

Mnemonic	Instruction	Type	Description
milemonie	instruction	type	Description
AND rd, rs1, rs2	AND	R	rd ← rs1 & rs2
OR rd, rs1, rs2	OR	R	rd ← rs1   rs2
XOR rd, rs1, rs2	XOR	R	rd ← rs1 ^ rs2
ANDI rd, rs1, imm12	AND immediate	I.	rd ← rs1 & imm12
ORI rd, rs1, imm12	OR immediate	I.	rd ← rs1   imm12
XORI rd, rs1, imm12	XOR immediate	I.	rd ← rs1 ^ imm12
SLL rd, rs1, rs2	Shift left logical	R	rd ← rs1 << rs2
SRL rd, rs1, rs2	Shift right logical	R	rd ← rs1 >> rs2
SRA rd, rs1, rs2	Shift right arithmetic	R	rd ← rs1 >> rs2
SLLI rd, rs1, shamt	Shift left logical immediate	I.	rd ← rs1 << shamt
SRLI rd, rs1, shamt	Shift right logical imm.	I.	rd ← rs1 >> shamt
SRAI rd, rs1, shamt	Shift right arithmetic immediate	I.	rd ← rs1 >> shamt

EIE2 Instruction Architectures & Compilers

Lecture 6 Slide 26

These instructions perform logical operations.

Mnemonic	Instruction	Туре	Description
LW rd, imm12(rs1)	Load word	I.	rd ← mem[rs1 + imm12]
LH rd, imm12(rs1)	Load halfword	I.	rd ← mem[rs1 + imm12]
LB rd, imm12(rs1)	Load byte	I.	rd ← mem[rs1 + imm12]
LWU rd, imm12(rs1)	Load word unsigned	T	rd ← mem[rs1 + imm12]
LHU rd, imm12(rs1)	Load halfword unsigned	I.	rd ← mem[rs1 + imm12]
LBU rd, imm12(rs1)	Load byte unsigned	Т	rd ← mem[rs1 + imm12]
SW rs2, imm12(rs1)	Store word	s	rs2(31:0) → mem[rs1 + imm12]
SH rs2, imm12(rs1)	Store halfword	s	rs2(15:0) → mem[rs1 + imm12]
SB rs2, imm12(rs1)	Store byte	s	rs2(7:0) → em[rs1 + imm12]

EIE2 Instruction Architectures & Compilers

Lecture 6 Slide 27

These instructions perform data memory read and write operations using pointer address in register and an immediate offset.

Mnemonic	Instruction	Туре	Description
BEQ rs1, rs2, imm12	Branch equal	SB	if rs1 == rs2 pc ← pc + imm12
BNE rs1, rs2, imm12	Branch not equal	SB	if rs1 != rs2 pc ← pc + imm12
BGE rs1, rs2, imm12	Branch greater than or equal	SB	if rs1 >= rs2 pc ← pc + imm12
BGEU rs1, rs2, imm12	Branch greater than or equal unsigned	SB	if rs1 >= rs2 pc ← pc + imm12
BLT rs1, rs2, imm12	Branch less than	SB	if rs1 < rs2 pc ← pc + imm12
BLTU rs1, rs2, imm12	Branch less than unsigned	SB	if rs1 < rs2 pc ← pc + imm12 << 1
JAL rd, imm20	Jump and link	IJ	rd ← pc + 4 pc ← pc + imm20
JALR rd, imm12(rs1)	Jump and link register	I	rd ← pc + 4 pc ← rs1 + imm12

EIE2 Instruction Architectures & Compilers

Lecture 6 Slide 28

These are the branch and jump instructions involving offset to the Program Counter.

# 28

Mnemonic	Instruction	Base instruction(s)	Mnemonic	Instruction	Base instruction(s)
LI rd, imm12	Load immediate (near)	ADDI rd, zero, imm12	BEQZ rs1, offset	Branch if rs1 = 0	BEQ rs1, zero, offset
LI rd, imm	Load immediate (far)	LUI rd, imm[31:12] ADDI rd, rd, imm[11:0]	BNEZ rs1, offset	Branch if rs1 ≠ 0	BNE rs1, zero, offset
LA rd, sym	Load address (far)	AUIPC rd, sym[31:12] ADDI rd, rd, sym[11:0]	BGEZ rs1, offset	Branch if rs1 ≥ 0	BGE rs1, zero, offset
MV rd, rs	Copy register	ADDI rd, rs, 0		Deres bill and and	
NOT rd, rs	One's complement	XORI rd, rs, -1	BLEZ rs1, offset	Branch if rs1 ≤ 0	BGE zero, rs1, offset
NEG rd, rs	Two's complement	SUB rd, zero, rs	BGTZ rs1, offset	Branch if rs1 > 0	BLT zero, rs1, offset
BGT rs1, rs2, offset	Branch if rs1 > rs2	BLT rs2, rs1, offset	J offset	Unconditional jump	JAL zero, offset
BLE rs1, rs2, offset	Branch if rs1 ≤ rs2	BGE rs2, rs1, offset	CALL offset12	Call subroutine (near)	JALR ra, ra, offset12
BGTU rs1, rs2, offset	Branch if rs1 > rs2 (unsigned)	BLTU rs2, rs1, offset	CALL offset	Call subroutine (far)	AUIPC ra, offset[31:12] JALR ra, ra, offset[11:0]
BLEU rs1, rs2, offset	Branch if rs1 ≤ rs2 (unsigned)	BGEU rs2, rs1, offset	RET	Return from subroutine	JALR zero, 0(ra)
	(unsigned)		NOP	No operation	ADDI Zero, Zero, 0

EIE2 Instruction Architectures & Compilers

Lecture 6 Slide 29

These are all the pseudo instructions accepted by the RISC-V assembler but are not really RISC-V instructions in the ISA. They are translated by the RISC-V assembler to one or more RISC-V instructions to make the program more readable.