


SOME ASSEMBLY REQUIRED



- What's up with immediates
- Let's Compile
- Let's Optimize



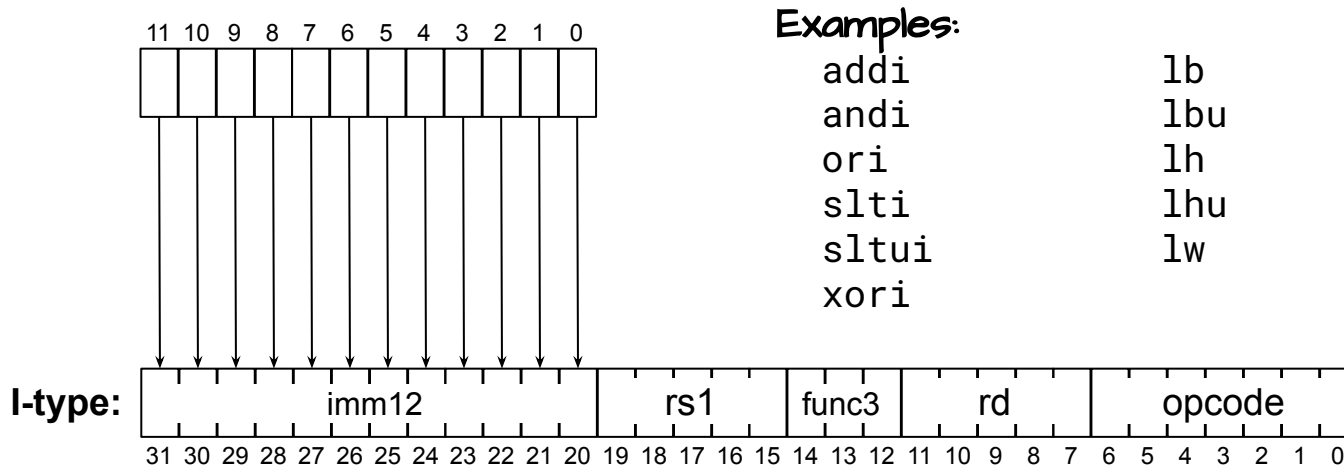
Midterm is one week from today!

Open book, open notes, open internet
(no communications however)



IMMEDIATE ENCODINGS


Many RISC-V instructions encode a constant value as part of the instruction. Unlike other ISAs, the encoding of constants is not uniform. The I-type instructions encode their immediate operand as follows:





IMMEDIATE ENCODINGS

Many RISC-V instructions encode a constant value as part of the instruction. Unlike other ISAs, the encoding of constants is not uniform. The I-type instructions encode their immediate operand as follows:




But there are exceptions!



LUI'S IMMEDIATE VALUE

The immediate-field encoding of `lui` and `auipc` also seems straightforward, but, as discussed last lecture, the immediate values might not be what you expect due to the sign-extension of its companion `addi` instruction.




Remember if bit 11 of the large constant is set, then you add 1 to the LUI/AUIPC immediate value.





IMMEDIATE OFFSETS FOR BRANCHES

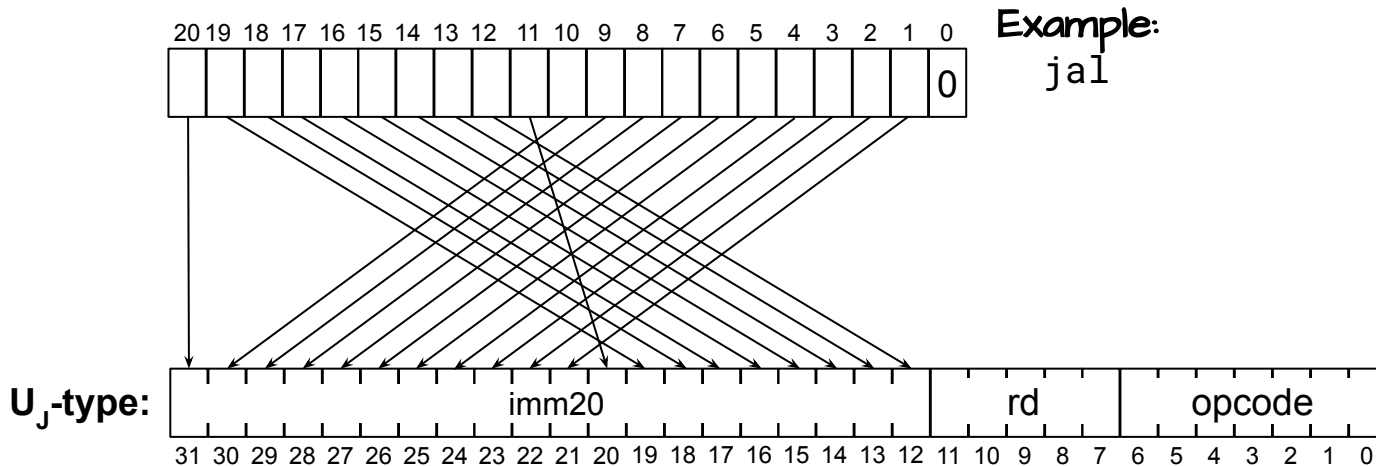
The offsets of RISC-V branch instructions use the same format as a stores, but the *offset encoding* is non-obvious. Again, there is no need for a rd register; only rs1 and rs2 are used. Branch offsets can only be even. In fact, all instructions must be aligned to half-word boundaries.





CRAZY JUMP OFFSETS

The offsets of RISC-V jump instructions use the same U-format as lui/auipc, but the *offset encoding* is even stranger. Jump offsets, like branch offsets can only be even.



Another bit-diddling
job for the assembler



WHY ASSEMBLE? COMPILE



There are a lot of details involved when mapping an assembly language program to bits in memory. But, this mapping can be automated. Compilers provide a means for hiding the details of assembly language. Lets, look at a simple recursive program:

```
int factorial(x) {
    if (x < 2)
        return x;
    return x * factorial(x-1);
}

int main() {
    return factorial(7);
}
```


THE MINIRISC-V C-COMPILER



```
UNC miniRISC-V C-compiler V 0.1

int factorial(x) {
    if (x < 2)
        return x;
    return x * factorial(x-1);
}

int main() {
    return factorial(7);
}
```

Include crt0.s Optimize

Goto

<https://csbi.unc.edu/mcmillan/index.py?run=rv>

This compiler does nly code generation. It bypasses the linking stage,

The assembly code that it generates can be pasted directly into the simulator

Cut and paste your C-code and press [Compile]

THE MINIRISC-V C-COMPILER



```
.align 4
.text
.globl factorial
.align 4
```

factorial:

```
addi x2,x2,-64
sw x8,60(x2)
addi x8,x2,48
sw x1,24(x2)
sw x27,28(x2)
```

```
mv x27,x12
li x30,2
bge x27,x30,L.2
addi x10,x27,0
jal x0,L.1
```

L.2:

```
addi x12,x27,-1
jal x1,factorial
addi x30,x10,0
mul x10,x27,x30
```

L.1:

```
lw x1,24(x2)
lw x27,28(x2)
lw x8,60(x2)
addi x2,x2,64
jalr x0,x1,0
```



A larger than necessary stack frame with a "frame pointer".

```
.globl main
.align 4
```

main:

```
addi x2,x2,-48
sw x8,44(x2)
addi x8,x2,32
sw x1,24(x2)
```

```
li x12,7
jal x1,factorial
addi x30,x10,0
```

L.4:

```
lw x1,24(x2)
lw x8,44(x2)
addi x2,x2,48
jalr x0,x1,0
```

```
.align 4
```

While all the assembly code generated is valid. It does not provide any support for initializing before startup. Rerun the compiler with "include crt0.s" selected.



THE C RUNTIME STARTUP CODE

```
reset: lui    sp,0xc0000    # initialize stack pointer
      addi   sp,sp,0xff0
      jal   ra,main      # call main
*halt: j     halt
```



#####

This small section of code initializes the stack pointer and calls the function "main".

This code is loaded into the .kernel memory section.

Note that the lui/addi involves a constant with bit 11 set.

RUN IT!



UNC miniRISC-V Architecture Simulator V 0.1

```
reset: lui    sp,0xc0000    # initialize stack pointer
      addi   sp,sp,0xff0
      jal   ra,main        # call main
*halt: j     halt

#####

      .align 4
      .text
      .globl factorial
      .align 4
factorial:
      addi x2,x2,-64
      sw  x8,60(x2)
      addi x8,x2,48
      sw  x1,24(x2)
      sw  x27,28(x2)
      mv  x27,x12
      li  x30,2
      bge x27,x30,L.2
      addi x10,x27,0
      jal x0,L.1

L.2:
      addi x12,x27,-1
      jal x1,factorial
      addi x30,x10,0
      mul x10,x27,x30

L.1:
      lw  x1,24(x2)
      lw  x27,28(x2)
```

Assemble [Reset] [Step] [Multistep] 10 [Run] [Memory Dump] [0x00000000] [Advanced]

Registers, Instruction Count = 131, Memory References = 177 pc: [0x0000000c]

x0/zero: [0x00000000]	x1/ra: [0x0000000c]	x2/sp: [0xBFFFFFF0]	x3/gp: [0x00000000]
x4/tp: [0x00000000]	x5/t0: [0x00000000]	x6/t1: [0x00000000]	x7/t2: [0x00000000]
x8/fp/s0: [0x00000000]	x9/s1: [0x00000000]	x10/a0: [0x000013B0]	x11/a1: [0x00000000]
x12/a2: [0x00000001]	x13/a3: [0x00000000]	x14/a4: [0x00000000]	x15/a5: [0x00000000]
x16/a6: [0x00000000]	x17/a7: [0x00000000]	x18/s2: [0x00000000]	x19/s3: [0x00000000]
x20/s4: [0x00000000]	x21/s5: [0x00000000]	x22/s6: [0x00000000]	x23/s7: [0x00000000]
x24/s8: [0x00000000]	x25/s9: [0x00000000]	x26/s10: [0x00000000]	x27/s11: [0x00000000]
x28/t3: [0x00000000]	x29/t4: [0x00000000]	x30/t5: [0x000013B0]	x31/t6: [0x00000000]

Address	Contents	Instruction
0x0001006C	0x02C12403	lw x8,44(x2)
0x00010070	0x03010113	addi x2,x2,48
0x00010074	0x00008067	jalr x0,x1,0
0x0000000c	0x0000006F	halt: j halt
0x00000010	0x00000000	[Invalid]
0x00000014	0x00000000	[Invalid]
0x00000018	0x00000000	[Invalid]

The code should assemble and run.

The simulator keeps track of the number of instructions executed. (131)

How could we count the number of calls to factorial?

ANOTHER EXAMPLE



Counts ones:

```
int countOnes(unsigned int x) {
    if (x == 0)
        return x;
    return (x & 1) + countOnes(x>>1);
}

int main() {
    return countOnes(1023);
}
```

What does this function do?

Compile and test it.

How would you improve it?

Speed?

Size?

NEXT TIME



We look into the hardware

