

Of course, the hardware doesn't really execute MIPS assembly language code.

The hardware can only store bits, and so the instructions it executes must be expressed in a suitable binary format.

We call the language made up of those instructions the *machine language*.

Different families of processors typically support different machine languages.

In the beginning, all programming was done in machine language... very ugly...

Assembly languages were created to make the programming process more human-centric.

Assembly language code is translated into machine language by an *assembler*.

Alas, there is no universal assembly language. In practice, assembly languages are tightly coupled with the underlying machine language and hardware.

Assembly provides convenient symbolic representation

- much easier than writing down numbers
- e.g., destination first

Machine language is the underlying reality

- e.g., destination is no longer first

Assembly can provide 'pseudoinstructions'

- e.g., “`move $t0, $t1`” exists only as an extension to assembly
- would be translated to “`add $t0, $t1, $zero`” by the assembler

When considering performance you should count real instructions

Examining the (basic) MIPS assembly instructions, we can easily identify three fundamentally different categories, according to the parameters they take:

- instructions that take 3 registers

```
add    $s0, $s1, $s2
```

```
or     $t1, $t0, $t1
```

- instructions that take 2 registers and an immediate (offset, number, address, etc.)

```
addi   $t7, $s5, 42
```

```
lw     $s3, 12($t4)
```

```
beq    $t1, $t3, Label01
```

- instructions that take an immediate

```
j      Label01
```

Overview of MIPS Machine Language

Simple instructions, all 32 bits wide

Very structured, no unnecessary baggage

Only three instruction formats:

R	op	rs	rt	rd	shamt	funct
I	op	rs	rt	16-bit immediate		
J	op	26-bit immediate				

R-format: basic arithmetical-logical instructions

I-format: load/store/conditional branch instructions

J-format: jump/unconditional branch instructions

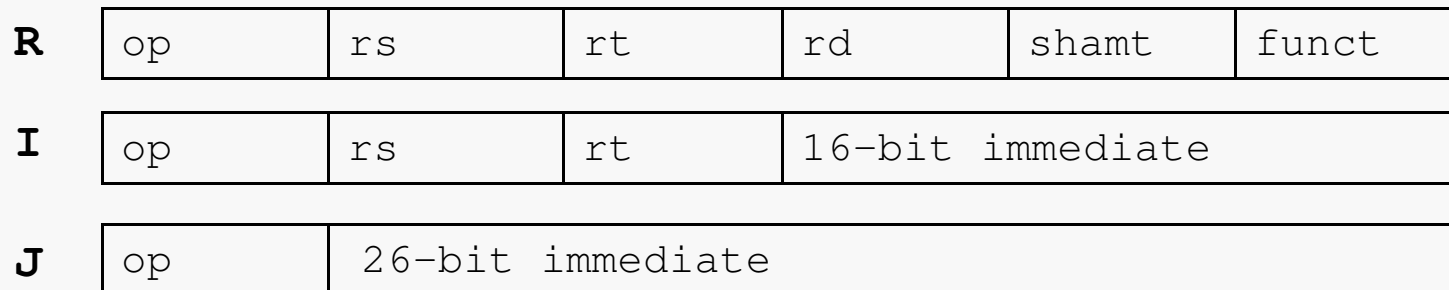
Overview of MIPS Machine Language

In MIPS32 Release 2, there are over 200 basic MIPS machine instructions.

In order to specify that many different instructions, we could use a field of 7 or more bits in every machine instruction.

But MIPS machine language uses a 6-bit *opcode* field and a variety of special cases.

For R-format instructions, the opcode field is set to 000000, and the last 6 bits specify exactly which arithmetic/logical instruction is to be performed.

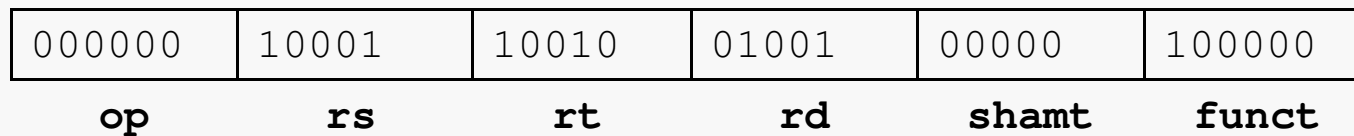


Instructions, like registers and words of data, are also 32 bits long

Example: `add $t1, $s1, $s2`

registers have numbers, $\$t1 = 9$, $\$s1 = 17$, $\$s2 = 18$

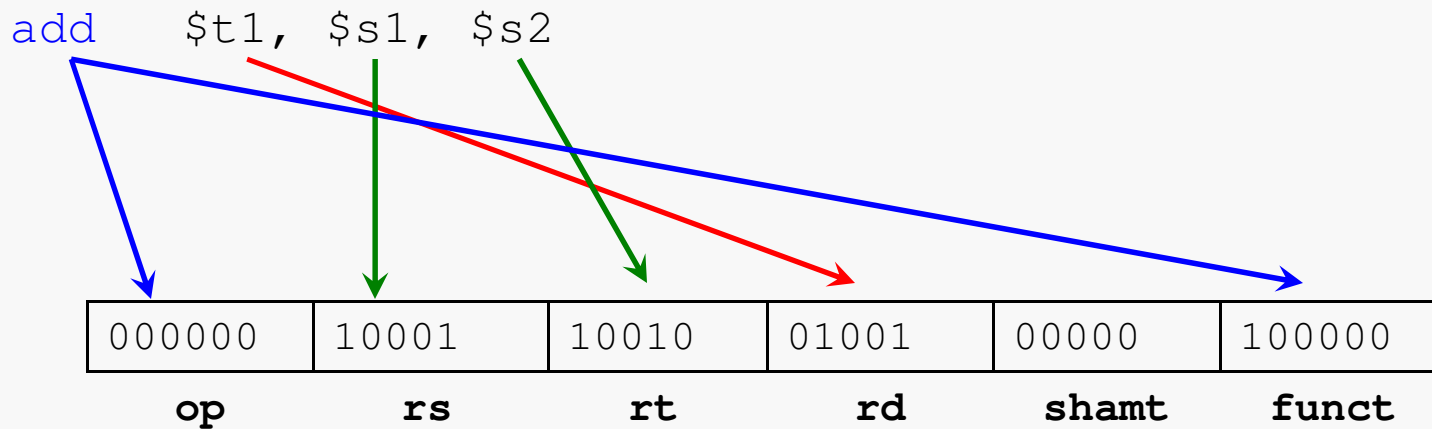
Machine language basic arithmetic/logic instruction format:



Can you guess what the field names stand for?

op	operation code (opcode)
rs	1st source register
rt	2nd source register
rd	destination register
shamt	shift amount
funct	opcode variant selector

Note how the assembly instruction maps into the machine representation:



The three register fields are each 5 bits wide. Why?

For arithmetic-logical instructions, both the **op** field and the **funct** field are used to specify the particular operation that is to be performed.

If you view memory contents, this would appear as 0x02324820.

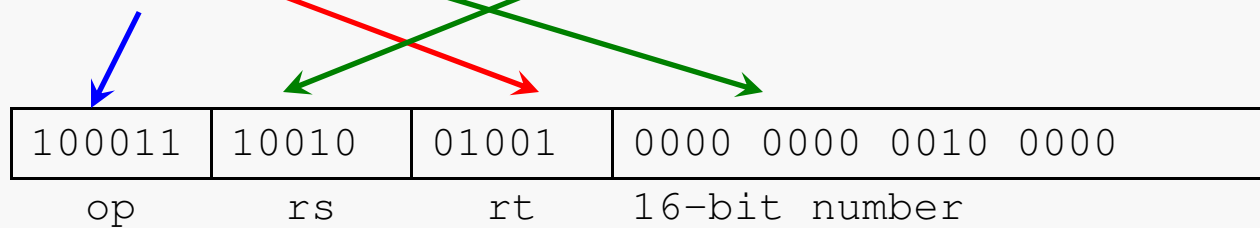
Consider the load-word and store-word instructions:

- what would the regularity principle have us do?
- new principle: Good design demands a compromise

We need a different type of machine language instruction format for these:

- I-type for data transfer instructions
- other format was R-type for register

Example: `lw $t0, 32($s2)`

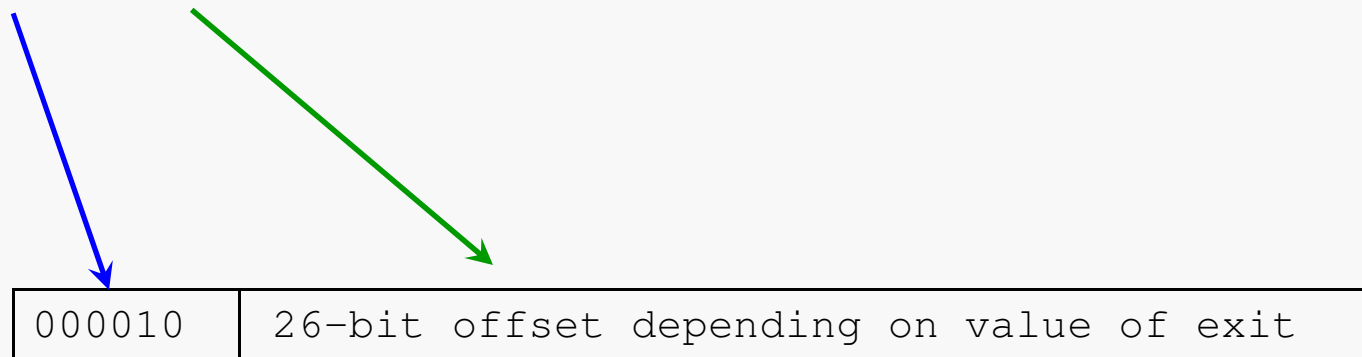


Where's the compromise?

Consider the jump instruction: `j Label01`

We need a different type of machine language instruction format for this as well.

Example: `j exit # assume exit is a statement label`



What will be involved in executing a machine language instruction?

Consider an I-type instruction, say a `lw` instruction:



The opcode bits must be analyzed to determine that the instruction is `lw`.

The contents of `$rs` must be fetched to the ALU and added to the immediate field to compute the appropriate address.

The contents at that address must be fetched from memory and written to register `$rt`.

We will need some sort of control logic that takes the 6-bit opcode and determines what specific machine instruction it corresponds to, and then triggers the necessary operations of other hardware units.

The control logic will be purely combinational.

Each hardware component, like the ALU will take certain control inputs; it will be the job of the decoder unit to set those control lines.

Note: for R-type instructions, this will be a little more involved.

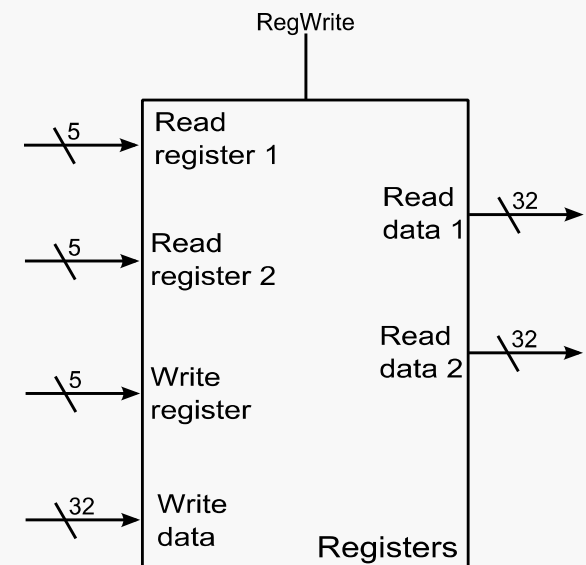
Fetching Register Operands

The registers must be organized in some way that makes it possible to conduct a transaction of the form "give me the contents of register number K".

Such a device is called a *register file*.

A register file has an input port that accepts a register number (a 5-bit port in our case) from which data will be read.

The register file processes that input and channels the contents of the corresponding register to an output port (a 32-bit port in our case).



The value fetched from the register must be combined with the 16-bit immediate value from the instruction.

This will require an adder circuit, but we will embed that in a more complex arithmetic-logic unit (ALU).

The ALU will accept two 32-bit values (operands) and accept control input bits that specify which operation is to be applied to those operands.

The ALU will supply the computed result to a 32-bit output port.

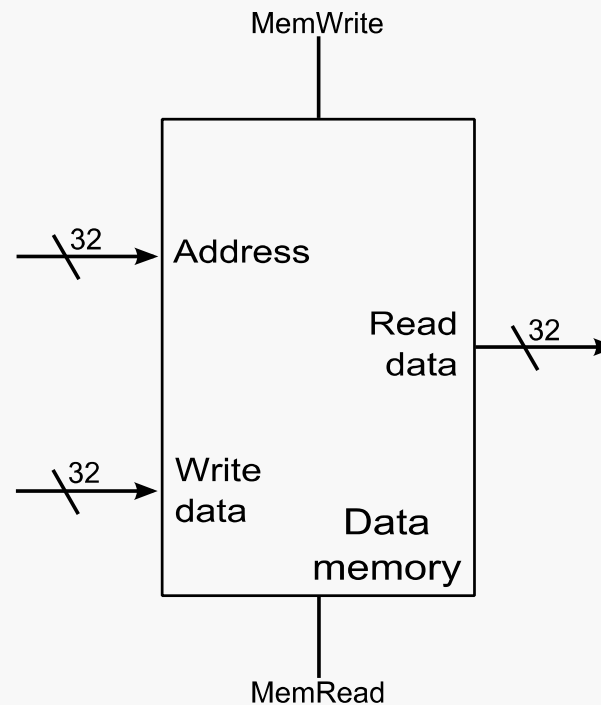
Note: the ALU will be designed to take 32-bit operands. That means that the 16-bit value taken from the instruction must be *widened* to 32-bits, and that must be done in such a way that the sign of the immediate value is preserved.

Fetching the Data

There must be a memory unit that has an input port that will accept a 32-bit address.

The memory unit must have a control input that specifies a value is to be read from that address (as opposed to being written to it).

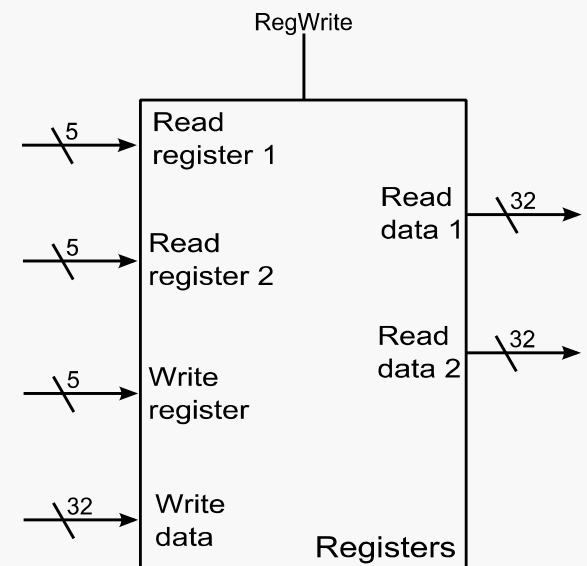
The memory unit must have a 32-bit output port to which the memory unit can channel the requested data value.



Loading the Data to a Register

The register file must have an input port that accepts a register number to which data will be written.

It will also need a control input to trigger the write operation.



A careful examination of the previous few slides will also reveal the need for a number of connections to channel data from one hardware component to another.

For example, there must be a way to move data from the register file's output port to the ALU's input ports.

In order to design a datapath for executing MIPS machine language instructions, we must identify all the hardware components we will need, and the necessary connections, and the necessary control logic and control lines...