

Slide 1

Administrivia

- (None? but I hope you all had a pleasant spring break?)

Slide 2

Designing a Processor — Overview

- Goal of Chapter 4 — sketch design of a (hardware) implementation of MIPS architecture in terms of some simple building blocks (AND and OR gates, inverters).
- Key components of the design (Figures 4.1 and 4.2):
 - Something to implement memory.
 - Something to implement instructions: “ALU” (arithmetic/logic unit).
 - Something to implement registers: “register file”.
 - Something to implement fetch/decode/execute cycle: “control logic”.

The first three together make up the “data path”. Analogy — it’s a puppet, with “control” pulling its strings.

Circuit Design — Overview

- AND and OR gates implement Boolean-algebra functions of the same names; inverter implements “not”.
- A word about notation: We’ll use the textbook’s notation for Boolean algebra, which alas is (probably?) different from what you used in CSCI 1323.

Slide 3

<i>CSCI 2321</i>	<i>CSCI 1323</i>
$a \cdot b$	$a \wedge b$
$a + b$	$a \vee b$
\bar{a}	a'

Implementing Logic Gates — Executive-Level Summary

- The ones and zeros of low-level software become two distinct voltages in hardware, and the logic of Boolean algebra is implemented using “switches” (things that connect an input to an output, or not, depending on the state of a control input).
- Currently these switches are (usually?) transistors. In widely-used “CMOS technology”, there are two types of switches, one that’s good if the input is “one” and one that’s good if the input is “zero”. These can be combined to implement logic. We looked earlier in the semester at a simple example (inverter) — link to description/explanation on “useful links” page.

Slide 4

Circuit Design — Overview Continued

Slide 5

- “Combinational logic” blocks implement Boolean functions/operations — map input(s) to output(s) without a notion of persistent state. (Think of these as “pure” functions that don’t change any variables but can have multiple outputs.)
- “Sequential logic” blocks also implement Boolean functions/operations but include a notion of persistent state. (Think of these as methods in object-oriented programming, which map input(s) to output(s) but also have access to member variables that can be read/written.)

Combinational Logic

Slide 6

- How to specify combinational logic block?
- One way — truth table with one line for each combination of inputs.
- Another way — Boolean-algebra expression(s) that define output(s) in terms of input(s).
- Example — 1-bit “adder”. Inputs are two digits to add and a carry-in, and outputs are sum and carry-out. So we would need a truth table with how many rows? how many columns? Or we could use Boolean expressions — how many? (See Figure B.5.3.)

Two-Level Logic

Slide 7

- Constructing logic blocks that implement arbitrary Boolean algebra expressions could take some thought.
- However, any Boolean-algebra expression can be represented in one of two forms — sum of products or product of sums. (Why? Think about truth-table representation.)

Two-Level Logic Implementations

Slide 8

- So we can define, for any combinational logic block, something that maps n inputs to m outputs by connecting an “array” of AND gates (one for each combination of inputs) to an “array” of OR gates (one for each output). (Example in Figure B.3.5.)
- Notice that representation in Figure B.3.5 could be changed to represent a different function by changing the positions of the dots — so generic term “programmable logic array” (PLA) makes sense?
- Another standardized way to represent combinational logic block is “ROM” (read-only memory) — for n inputs and m outputs we’d need 2^n entries each consisting of m bits.
- For either of these the process of turning a truth table into implementation can be automated(!).

Slide 9

“Managing Complexity”

- Worth noting that, as in programming, the discussion will make extensive use of layers of abstraction to build complex things from simple things.
- Just as in programming it's common to define library functions that implement frequently-used operation, we can define some not-so-basic blocks, such as decoders and multiplexors. (See discussion in B.3, especially Figures B.3.2 and B.3.1.)

Slide 10

“Don't Care” Inputs/Outputs

- For not-so-small numbers of inputs a full truth table can be big, so it's worthwhile to think about whether there's something simpler that gets the same effect.
- One way to do this — exploit “don't care”s. Input “don't care” arises when both values for an input (in combination with other inputs) give same result. Output “don't care” arises when we aren't interested in output for some combination of inputs (maybe it can never occur?). Textbook shows how to use this idea to produce a shorter truth table.
- Exploiting the shorter table, and in general minimizing the complexity of the combinational logic block, can be done manually (“Karnaugh maps”) or automatically (various design tools).

Arrays of Logic Elements

Slide 11

- Descriptions so far (except for decoder) have been in terms of single-bit inputs. But often we want to work on larger collections (e.g., the 32 bits of a register).
- To do this, we (usually?) can build an “array” of identical logic blocks.
- If inputs/outputs are not in some way connected, can just indicate that input/output values are more than one bit (“bus”). Examples — bitwise AND of 32-bit values, Figure B.3.6.
- If inputs/outputs *are* connected, idea still works but picture must indicate connections. Example — addition of 32-bit values using 32 single-bit “adder” blocks, each with three inputs (two operands and carry-in) and two outputs (value and carry-out).

Design of an ALU

Slide 12

- One of the things we need for a MIPS implementation is something that can do the arithmetic and logic operations in the MIPS instruction set.
- Inputs to operations are typically two 32-bit values. Some operations can be done by operating on all bits in exactly the same way and independently (e.g., and). Others can be done by operating on all bits in the same way but with dependencies among bits (e.g., add). So we will design a “1-bit ALU” and then figure out how to connect 32 of them to make the full 32-bit logic block.

Slide 13

1-Bit ALU

- Figures B.5.1 through B.5.6 show how we can build up something that performs `and`, `or`, and `add` on 1-bit values (plus carry-in and carry-out values for `add`).
- Result (B.5.6) is a logic block with inputs
 - two 1-bit operands
 - 2-bit “which operation?”
 - 1-bit carry-inand outputs
 - 1-bit result
 - 1-bit carry-out

Slide 14

32-Bit ALU from 1-Bit ALUs

- Now we want to connect 32 of these 1-bit ALUs to make a 32-bit ALU.
- Figure B.5.7 shows how:
 - Connect operand inputs of each 1-bit ALU to individual bits of 32-bit operand, and similarly for 32-bit result.
 - Connect “which operation?” input (common to all) to “which operation?” input of each 1-bit ALU.

Slide 15

32-Bit ALU from 1-Bit ALUs, Continued

- We said when we first talked about two's complement notation that it was attractive because once you build something that can add, you can easily extend it to something that can subtract, right?
- Conceptually, we can compute $a - b$ by adding a to $-b$, and we can compute $-b$ by reversing all the bits of b and adding one — which is just what's shown in Figure B.5.8! which is Figure B.5.7 plus one more input, which:
 - if 0, makes the initial carry-in 0 and uses b as is.
 - if 1, makes the initial carry-in 1 and flips bits of b .
- We can apply a similar idea (adding an input that lets us use a as is or “flipped”) to implement `nor` (Figure B.5.9).

Slide 16

32-Bit ALU from 1-Bit ALUs, Continued

- Figures B.5.10 and B.5.11 and accompanying text show how to extend the design to implement `slt` and also an overflow detector. Executive-level summary: Calculate $a - b$ and use high-order bit of result of that operation to set low-order bit of result.
- Result is something we can use to do pretty much all of the arithmetic and logic operations of the MIPS ISA. Exceptions are shifts (but those don't seem like they'd be too hard) and multiplication/division (which do, so skip for now).
Notice also that getting valid output values may take a while for some operations, such as addition — values “flow” through the circuit. Designers of real hardware use clever tricks to speed up addition, such as the one(s) described in B.6. Read if interested!

Memory Elements

Slide 17

- So now we (sort of) know how to design logic blocks that use switches/gates to compute output bits from input bits.
- But where do those input bits come from, and where do the output bits go? “state elements” — things that can save values.
- (Keep in mind that the goal here is to get a sense of how you can build something that stores a value out of gates/switches. Details are (I think!) very interesting but can to some extent be skimmed.)

Minute Essay

Slide 18

- We sketched a somewhat-simple design for a 32-bit ALU. We could make a 64-bit ALU in much the same way. Comparing the two in terms of how long it would take to do each of the discussed operations, which would you guess to be faster (if either)?
- Does the answer to the previous question depend on which instruction is being executed?

Minute Essay Answer

- The 64-bit ALU will be slower for some operations (such as `add`), since “values” have “flow” through 64 1-bit ALUs rather than 32.
(However, as one student pointed out, if the ALU is doing all the operations anyway even though only one is being used, in some sense they do all take the same amount of time.)

Slide 19