MITOCW | MIT6_004S17_03-02-01_300k

Let's review our wish list for the characteristics of a combinational device.

In the previous lecture we worked hard to develop a voltage-based representation for information that could tolerate some amount error as the information flowed through a system of processing elements.

We specified four signaling thresholds: V_OL and V_OH set the upper and lower bounds on voltages used to represent 0 and 1 respectively at the outputs of a combinational device.

V_IL and V_IH served a similar role for interpreting the voltages at the inputs of a combinational device.

We specified that V_OL be strictly less than V_IL, and termed the difference between these two low thresholds as the low noise margin, the amount of noise that could be added to an output signal and still have the signal interpreted correctly at any connected inputs.

For the same reasons we specified that V_IH be strictly less than V_OH.

We saw the implications of including noise margins when we looked at the voltage transfer characteristic - a plot of V_OUT vs. V_IN - for a combinational device.

Since a combinational device must, in the steady state, produce a valid output voltage given a valid input voltage, we can identify forbidden regions in the VTC, which for valid input voltages identify regions of invalid output voltages.

The VTC for a legal combinational device could not have any points that fall within these regions.

The center region, bounded by the four threshold voltages, is narrower than it is high and so any legal VTC has to a have region where its gain is greater than 1 and the overall VTC has to be non-linear.

The VTC shown here is that for a combinational device that serves as an inverter.

If we're fortunate to be using a circuit technology that provides high gain and has output voltages close the ground and the power supply voltage, we can push V_OL and V_OH outward towards the power supply rails, and push V_IL and V_IH inward, with the happy consequence of increasing the noise margins - always a good thing!

Remembering back to the beginning of Lecture 2, we'll be wanting billions of devices in our digital systems, so each device will have to be quite inexpensive and small.

In today's mobile world, the ability to run our systems on battery power for long periods of time means that we'll want to have our systems dissipate as little power as possible.

Of course, manipulating information will necessitate changing voltages within the system and that will cost us some amount of power.

But if our system is idle and no internal voltages are changing, we'd like for our system to have zero power dissipation.

Finally, we'll want to be able to implement systems with useful functionality and so need to develop a catalog of the logic computations we want to perform.

Quite remarkably, there is a circuit technology that will make our wishes come true!

That technology is the subject of this lecture.

The star of our show is the metal-oxide-semiconductor field-effect transistor, or MOSFET for short.

Here's a 3D drawing showing a cross-section of a MOSFET, which is constructed from a complicated sandwich of electrical materials as part of an integrated circuit, so called because the individual devices in an integrated circuit are manufactured en-masse during a series of manufacturing steps.

In modern technologies the dimensions of the block shown here are a few 10's of nanometers on a side - that's 1/1000 of the thickness of a thin human hair.

This dimension is so small that MOSFETs can't be viewed using an ordinary optical microscope, whose resolution is limited by the wavelength of visible light, which is 400 to 750nm.

For many years, engineers have been able to shrink the device dimensions by a factor of 2 every 24 months or so, an observation known as "Moore's Law" after Gordon Moore, one of the founders of Intel, who first remarked on this trend in 1965.

Each 50% shrink in dimensions enables integrated circuit (IC) manufacturers to build four times as many devices in the same area as before, and, as we'll see, the devices themselves get faster too!

An integrated circuit in 1975 might have had 2500 devices; today we're able to build ICs with two to three billion devices.

Here's a quick tour of what we see in the diagram.

The substrate upon which the IC is built is a thin wafer of silicon crystal which has had impurities added to make it conductive.

In this case the impurity was an acceptor atom like Boron, and we characterize the doped silicon as a p-type semiconductor.

The IC will include an electrical contact to the p-type substrate, called the "bulk" terminal, so we can control its voltage.

When want to provide electrical insulation between conducting materials, we'll use a layer of silicon dioxide (SiO_2).

Normally the thickness of the insulator isn't terribly important, except for when it's used to isolate the gate of the transistor (shown here in red) from the substrate.

The insulating layer in that region is very thin so that the electrical field from charges on the gate conductor can easily affect the substrate.

The gate terminal of the transistor is a conductor, in this case, polycrystalline silicon.

The gate, the thin oxide insulating layer, and the p-type substrate form a capacitor, where changing the voltage on the gate will cause electrical changes in the p-type substrate directly under the gate.

In early manufacturing processes the gate terminal was made of metal, and the term "metal-oxide-semiconductor" (MOS) is referring to this particular structure.

After the gate terminal is in place, donor atoms such as Phosphorous are implanted into the p-type substrate in two rectangular regions on either side of the gate.

This changes those regions to an n-type semiconductor, which become the final two terminals of the MOSFET, called the source and the drain.

Note that source and drain are usually physically identical and are distinguished by the role they play during the operation of the device, our next topic.

As we'll see in the next slide, the MOSFET functions as a voltage-controlled switch connecting the source and drain terminals of the device.

When the switch is conducting, current will flow from the drain to the source through the conducting channel formed as the second plate of the gate capacitor.

The MOSFET has two critical dimensions: its length L, which measures the distance the current must cross as it flows from drain to source, and its width W, which determines how much channel is available to conduct current.

The the current, termed I_DS, flowing across the switch is proportional to the ratio of the channel's width to its length.

Typically, IC designers make the length as short as possible.

When a news article refers to a "14nm process," the 14nm refers to the smallest allowable value for the channel length.

And designers choose the channel width to set the desired amount of current flow.

If I_DS is large, voltage transitions on the source and drain nodes will be quick, at the cost of a physically larger device.

To summarize: the MOSFET has four electrical terminals: bulk, gate, source, and drain.

Two of the device dimensions are under the control of the designer: the channel length, usually chosen to be as small as possible, and the channel width chosen to set the current flow to the desired value.

It's a solid-state switch - there are no moving parts - and the switch operation is controlled by electrical fields determined by the relative voltages of the four terminals.