

Flip-flops

- The second part of CS231 focuses on **sequential circuits**, in which we add memory to our circuits.
- The schedule will be very similar the first third of the class.
 - We first show how primitive memory units are built.
 - Then we talk about analysis and design of sequential circuits.
 - We also present common devices like registers and counters.
- In the final third of the summer, we'll add memory to our ALU, to make a complete processor!




Combinational circuits



- So far we've only worked with **combinational circuits**, where applying the same inputs always produces the same outputs.
 - This corresponds to a mathematical function, where every input has a single, unique output.
 - In programming terminology, combinational circuits are similar to “functional programs” that do not contain variables and assignments.
- Such circuits are comparatively easy to design and analyze.


Sequential circuits



- In contrast, the outputs of a **sequential circuit** depend on not only the inputs, but also the **state**, or the current contents of some memory.
- This makes things more difficult to understand since the same inputs can yield *different* outputs, depending on what's stored in memory.
- The memory contents can also change as the circuit runs, so the *order* in which things occur makes a difference.


Examples of sequential devices

- Many real-life devices are sequential in nature.
 - Combination locks open if you enter numbers in the right order.
 - Elevators move up or down and open or close in response to buttons that are pressed on different floors and in the elevator itself.
 - Traffic lights change from red to green depending on whether a car is waiting at the intersection.
- More importantly for us, computers are sequential! For instance, key presses and mouse clicks have different effects based on which program is loaded into memory, and the state of that program.





What exactly is memory?

- A memory should support at least three operations.
 1. It should be able to hold a value.
 2. You should be able to *read* the value that is saved.
 3. You should be able to *change* that value.
- We'll start with the simplest case, a one-bit memory.
 1. It should be able to hold a single bit, 0 or 1.
 2. You should be able to read the bit that is saved.
 3. You should be able to change the bit.
 - You can **set** the bit to 1
 - You can **reset** or **clear** the bit to 0.



The basic idea of storage


- How can a circuit remember anything, when it's just a bunch of gates that produce outputs according to the inputs?
- The idea is to make a loop in a circuit, so the outputs are *also* inputs.
- Here is one initial attempt, shown with two equivalent layouts.



- Does this satisfy the properties of memory?
 - These circuits “remember” Q since its value never changes. Similarly, Q' never changes either.
 - We can “read” Q by sending it to another gate or device, as usual.
 - But we can't *change* Q ! There are no external inputs here, so we can't control whether $Q=1$ or $Q=0$.

A really confusing circuit

- Let's use NOR gates instead of inverters. The **SR latch** here has two inputs S and R, which will let us control the outputs Q and Q'.



- Q and Q' feed back into the circuit, so they're not only outputs, they're also inputs!
- To figure out how Q and Q' change, we must look at not only the inputs S and R, but also the *current* values of Q and Q'.

$$Q_{\text{next}} = (R + Q'_{\text{current}})'$$

$$Q'_{\text{next}} = (S + Q_{\text{current}})'$$


- Let's see how different input values for S and R affect this thing.

Storing a value: SR = 00

- What if $S = 0$ and $R = 0$?
- The equations on the right reduce to:

$$Q_{\text{next}} = (0 + Q'_{\text{current}})' = Q_{\text{current}}$$
$$Q'_{\text{next}} = (0 + Q_{\text{current}})' = Q'_{\text{current}}$$

- So when $SR = 00$, then $Q_{\text{next}} = Q_{\text{current}}$.
- This is exactly what we need to **store** values in the latch.



$$Q_{\text{next}} = (R + Q'_{\text{current}})'$$
$$Q'_{\text{next}} = (S + Q_{\text{current}})'$$

Setting the latch: SR = 10


- What if $S = 1$ and $R = 0$?
- Since $S = 1$, Q'_{next} is 0, *regardless* of Q_{current} .

$$Q'_{\text{next}} = (1 + Q_{\text{current}})' = 0$$

- Then this new value of Q' goes into the top NOR gate, along with $R = 0$.

$$Q_{\text{next}} = (0 + 0)' = 1$$


- So when $SR = 10$, then $Q'_{\text{next}} = 0$ and $Q_{\text{next}} = 1$. This is how you **set** the latch to 1; the S input stands for “set.”
- Notice it can take up to two steps (two gate delays) from the time S becomes 1 to the time Q_{next} becomes 1.
- But once Q_{next} becomes 1, the outputs will stop changing. This is a **stable state**.



$$Q_{\text{next}} = (R + Q'_{\text{current}})'$$
$$Q'_{\text{next}} = (S + Q_{\text{current}})'$$

Latch delays

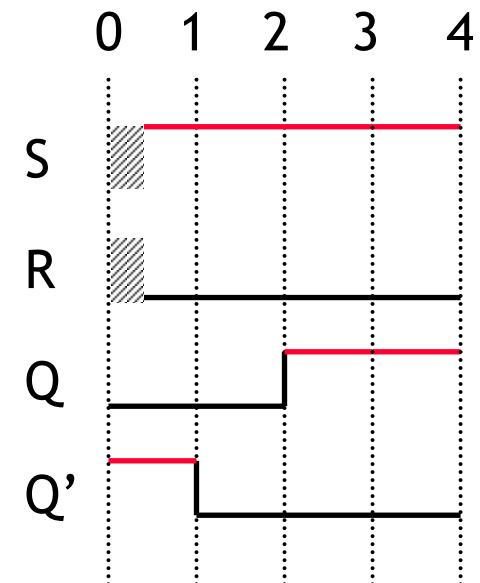
- Timing diagrams are especially useful for seeing how sequential circuits work.
- Here is a diagram which shows an example of how our latch outputs change with inputs $SR = 10$.



- Let's say that $Q = 0$ and $Q' = 1$ initially.
- Since $S = 1$, Q' changes from 1 to 0 after one NOR-gate delay (marked with vertical lines in the timing diagram).
- This change in Q' , along with $R = 0$, causes Q to become 1 after another gate delay.
- The latch then stabilizes until S or R change again.

$$Q_{\text{next}} = (R + Q'_{\text{current}})'$$

$$Q'_{\text{next}} = (S + Q_{\text{current}})'$$



Resetting the latch: SR = 01


- What if $S = 0$ and $R = 1$?
- Since $R = 1$, Q_{next} is 0, *regardless* of Q_{current} .

$$Q_{\text{next}} = (1 + Q'_{\text{current}})' = 0$$

- Then this new value of Q goes into the bottom NOR gate, where $S = 0$.

$$Q'_{\text{next}} = (0 + 0)' = 1$$

- So when $SR = 01$, then $Q_{\text{next}} = 0$ and $Q'_{\text{next}} = 1$. This is how you **reset**, or **clear**, the latch to 0; the R input stands for “reset.”
- Again, it can take two gate delays before a change in R propagates to the output Q'_{next} .



$$Q_{\text{next}} = (R + Q'_{\text{current}})'$$
$$Q'_{\text{next}} = (S + Q_{\text{current}})'$$

SR latches are memories!

- This **characteristic table** shows that our latch provides everything we need in a memory: we can set it, reset it, or keep the current value.
- The output Q represents the data stored in the latch. It is also called the **state** of the latch.
- We can expand the table above into a **state table**, which explicitly shows that the *next* values of Q and Q' depend on their *current* values, as well as on the inputs S and R.

S	R	Q
0	0	No change
0	1	0 (reset)
1	0	1 (set)

Inputs		Current		Next	
S	R	Q	Q'	Q	Q'
0	0	0	1	0	1
0	0	1	0	1	0
0	1	0	1	0	1
0	1	1	0	0	1
1	0	0	1	1	0
1	0	1	0	1	0

SR latches are sequential!

- Note that for $SR = 00$, the next value of Q could be *either* 0 or 1, depending on the current value of Q .
- So the same inputs can produce different outputs, depending on whether the latch is currently set or reset.
- This is different from the combinational circuits that we've seen so far, where the same inputs always generate the same outputs.

S	R	Q
0	0	No change
0	1	0 (reset)
1	0	1 (set)

Inputs		Current		Next	
S	R	Q	Q'	Q	Q'
0	0	0	1	0	1
0	0	1	0	1	0
0	1	0	1	0	1
0	1	1	0	0	1
1	0	0	1	1	0
1	0	1	0	1	0

What about SR = 11?

- Both Q_{next} and Q'_{next} would become 0, which contradicts the assumption that Q and Q' are always complements.
- Another problem is what happens if we then make $S = 0$ and $R = 0$ together.

$$Q_{\text{next}} = (0 + 0)' = 1$$

$$Q'_{\text{next}} = (0 + 0)' = 1$$



- But these new values go back into the NOR gates, and we then get $Q = Q' = 0$ again.

$$Q_{\text{next}} = (0 + 1)' = 0$$

$$Q'_{\text{next}} = (0 + 1)' = 0$$


- So the circuit enters an infinite loop, where Q and Q' cycle between 0 and 1 forever.
- Don't set $SR = 11!$

$$Q_{\text{next}} = (R + Q'_{\text{current}})'$$
$$Q'_{\text{next}} = (S + Q_{\text{current}})'$$



S'R' latch

- There are several other variations of the basic latch.
- You can use NAND instead of NOR gates to get a **S'R' latch**.

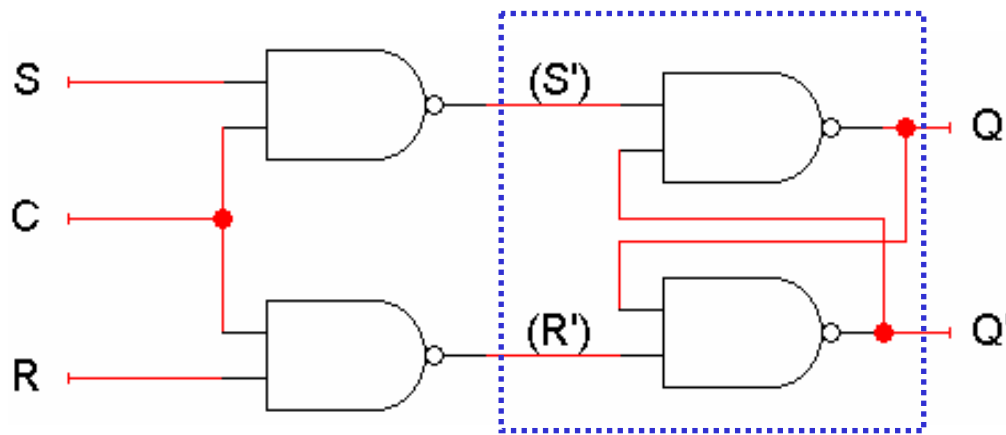


S'	R'	Q
1	1	No change
1	0	0 (reset)
0	1	1 (set)
0	0	Avoid!

- This is just like an SR latch but with inverted inputs, as you can see from the table.

An SR latch with a control input

- Here is an SR latch with a **control input** C, which acts like an enable.

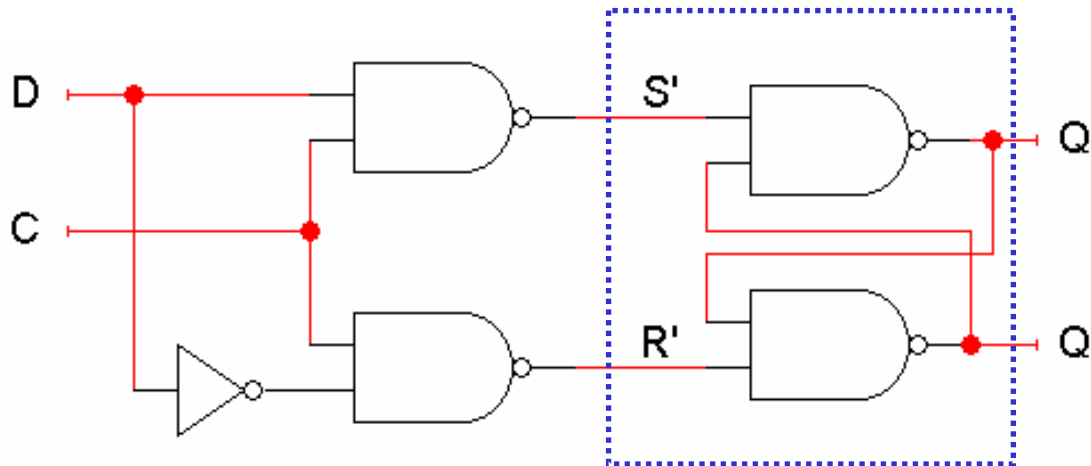


C	S	R	S'	R'	Q
0	x	x	1	1	No change
1	0	0	1	1	No change
1	0	1	1	0	0 (reset)
1	1	0	0	1	1 (set)
1	1	1	0	0	Evil!

- Notice the hierarchical design!
 - The dotted blue box contains the S'R' latch from the previous slide.
 - The additional NAND gates are simply used to generate appropriate inputs for the S'R' latch.
- We'll see more of the control input later today.

D latch

- A **D latch** is also based on an S'R' latch. The additional gates generate the S' and R' signals, based on inputs D ("data") and C ("control").
 - When $C = 0$, S' and R' are both 1, so Q does not change.
 - When $C = 1$, the latch output Q will equal the input D.



C	D	Q
0	x	No change
1	0	0
1	1	1


- There are two main advantages of a D latch.
 - No more messing with one input for set and another input for reset!
 - This latch has no "bad" input combinations to avoid. Any of the four possible assignments to C and D are valid.

Using latches in real life

- We can use some D latches as a memory for an ALU. The latches should normally be disabled, so unwanted data doesn't accidentally get stored.



- Let's say the latches initially hold 0000, and we want to increment them. The ALU can first read and increment the current latch contents.




Writing to the latches

- The latches should be enabled only *after* the ALU finishes the increment operation, so the updated value can be stored.



- The latch must be quickly disabled again, *before* the ALU has a chance to read the new value 0001 and produce a new result 0010.



Two main issues

- So to use latches within a circuit, we have to remember two things.
 - Keep the latches disabled until new values are ready to be stored.
 - Enable the latches just long enough for the update to occur.
- There are two main issues we need to address.

How do we know exactly when the new values are ready?

We'll add another signal to our circuit. When this new value changes to 1, the latches will know that the ALU computation completed and data is ready to be stored.

How can we enable and then quickly disable the latches?

This can be done by combining latches together in a special way, to form what are called **flip-flops**.

Clocks and synchronization

- A **clock** is a special device that continuously outputs 0s and 1s.
 - The time it takes the clock to change from 1 to 0 and back to 1 is called the **clock period**, or **clock cycle time**.
 - The **clock frequency** is the inverse of the clock period. The unit of measurement for frequency is the **hertz**.

clock period



- Clocks are often used to synchronize circuits.
 - They generate a repeating, predictable pattern of 0s and 1s that can trigger certain events in a circuit, such as writing to a latch.
 - If several circuits share a common clock signal, they can coordinate their actions with respect to one another.
- This is similar to how humans use real clocks for synchronization.

Synchronizing our example

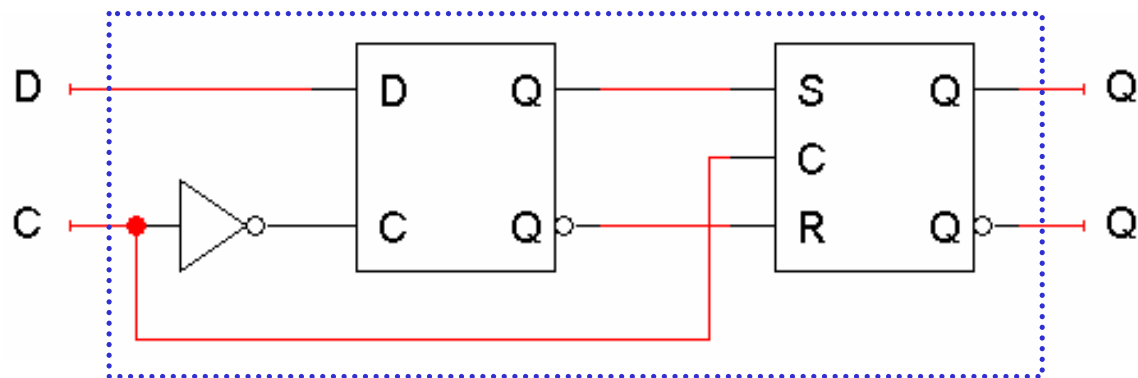
- We can use a clock to synchronize our latches with the ALU.
 - The clock signal is connected to the latch control input C.
 - The clock controls the latches. When it becomes 1, the latches will be enabled for writing.



- The clock period must be set appropriately for the ALU.
 - It should not be too short. Otherwise, the latches will start writing before the ALU operation has finished.
 - It should not be too long either. Otherwise the ALU may produce a new result that will get stored accidentally, as we saw before.
- The faster the ALU runs, the shorter the clock period can be.


Flip-flops

- The second issue was how to enable a latch for just an instant.
- Here is the internal structure of a **D-type flip-flop**.
 - The flip-flop inputs are C and D, and the outputs are Q and Q'.
 - The D latch on the left is the **master**, while the SR latch on the right is called the **slave**.



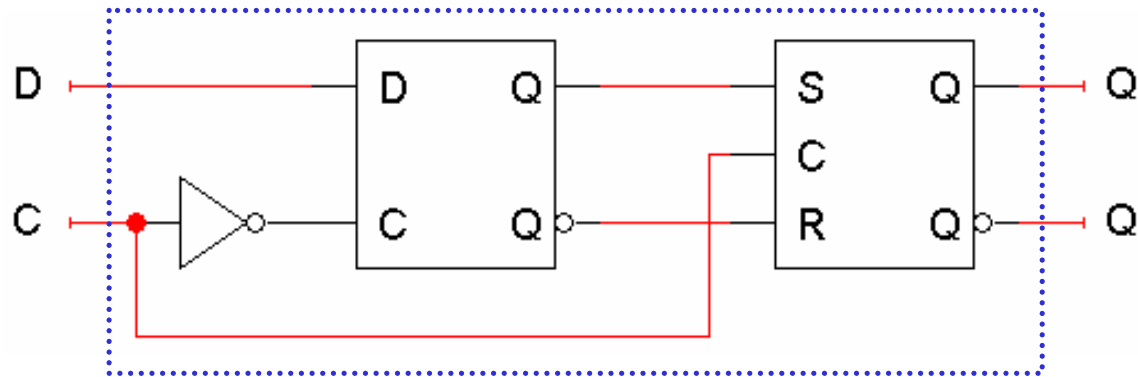
- Note the layout here.
 - The flip-flop input D is connected directly to the master latch.
 - The master latch output goes to the slave.
 - The flip-flop outputs come directly from the slave latch.

D flip-flops when C=0



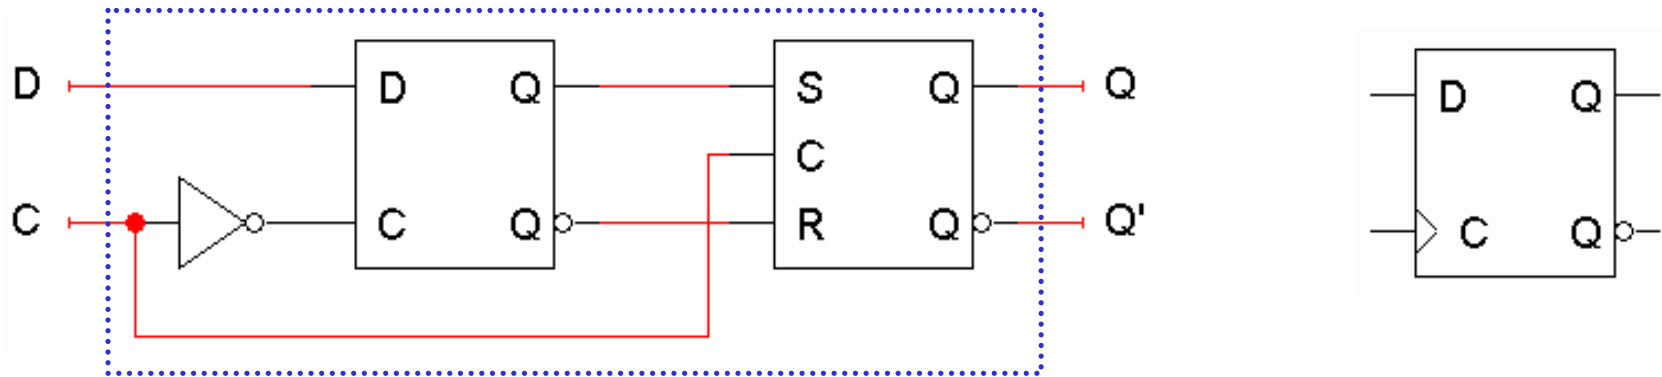
- The D flip-flop's control input C enables either the master D latch or the slave SR latch, but *not* both.
- What happens when C = 0?
 - The master latch is enabled, and it tracks the flip-flop input D. When D changes, the master's output changes too.
 - The slave is disabled, so the D latch output has no effect on it. Thus, the slave just maintains the flip-flop's current state.

D flip-flops when C=1



- Several things happen *as soon as* C changes from 0 to 1.
 - The master is disabled. Its output will be the *last* D input value seen, just before C became 1. Any subsequent changes to the D input will have no effect on the master latch while C = 1.
 - On the other hand, the slave is enabled. Its output changes to reflect the master's state, which again is the D input value from right when C became 1.

Positive edge triggering

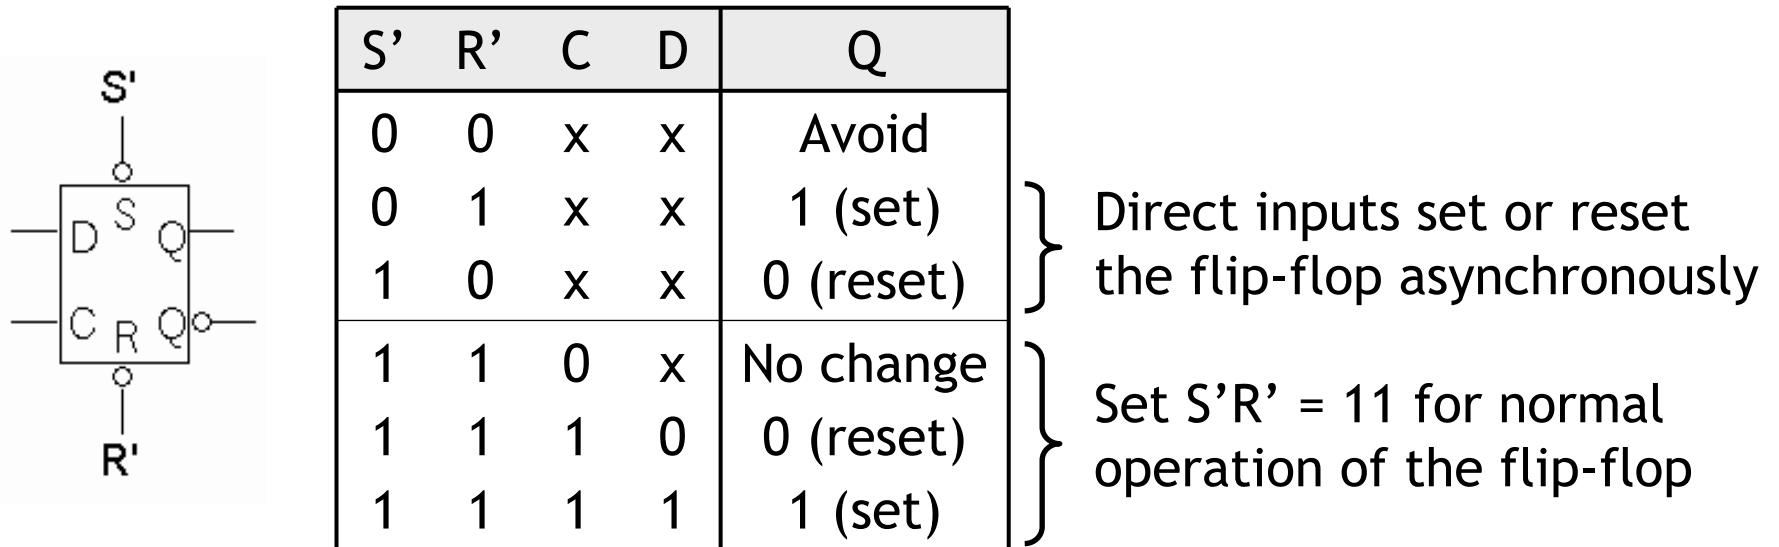


- This is called a **positive edge-triggered** flip-flop.
 - The flip-flop output Q changes *only* after the positive edge of C .
 - The change is based on the flip-flop input values that were present right at the positive edge of the clock signal.
- The D flip-flop's behavior is similar to that of a D latch, except for the positive edge-triggered nature, which is not explicit in this table.

C	D	Q
0	x	No change
1	0	0 (reset)
1	1	1 (set)

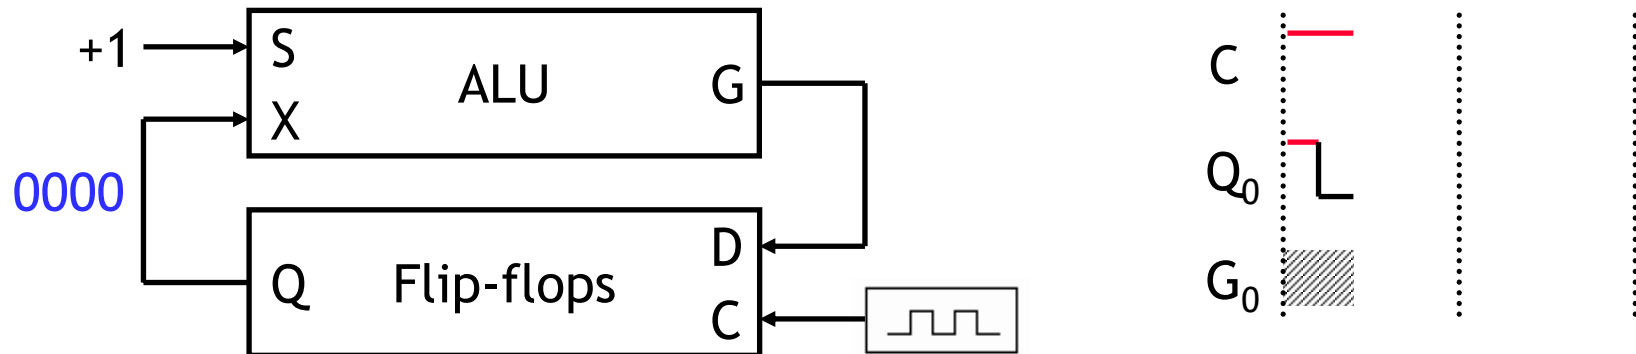
Direct inputs

- One last thing to worry about... what is the starting value of Q?
- We could set the initial value synchronously, at the next positive clock edge, but this actually makes circuit design more difficult.
- Most flip-flops provide **direct inputs**, or **asynchronous inputs**, that let you immediately set or clear the state, regardless of the clock input.
 - You would “reset” the circuit once, to initialize the flip-flops.
 - The circuit would then begin its regular, synchronous operation.
- Here is a LogicWorks D flip-flop with active-low direct inputs.

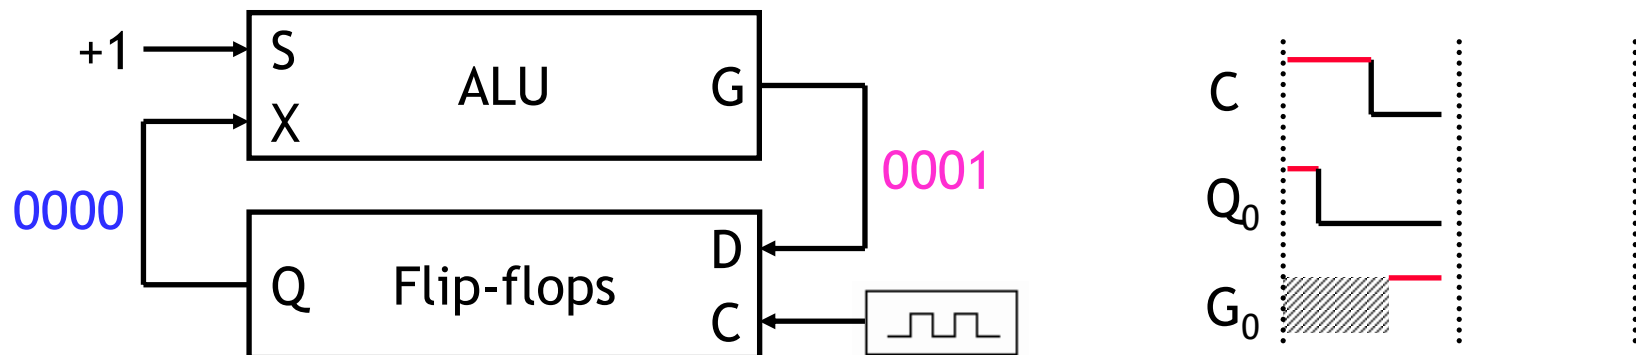


Our example with flip-flops

- We can use the flip-flops' direct inputs to initialize them to 0000.

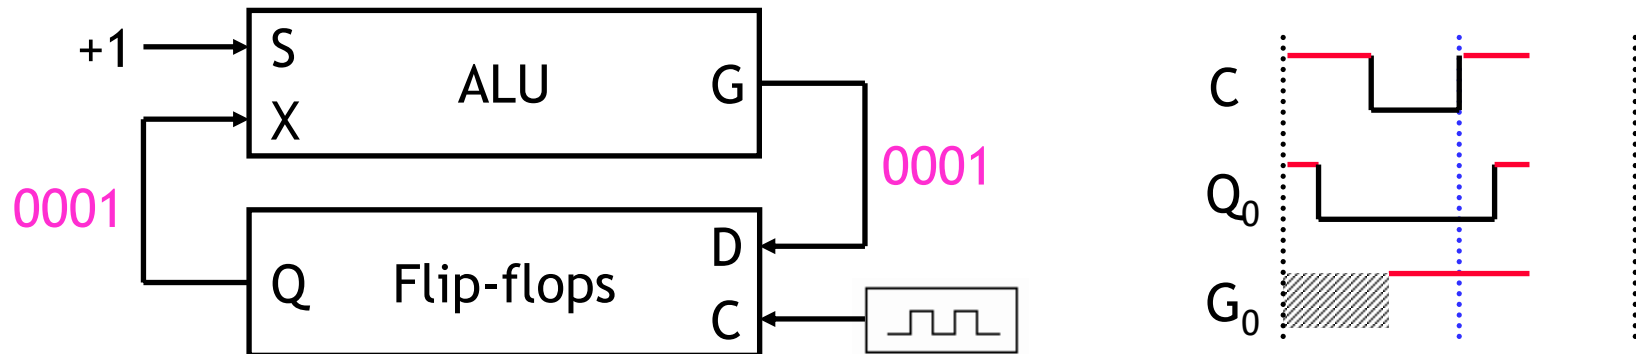


- During the clock cycle, the ALU outputs 0001, but this does not affect the flip-flops yet.

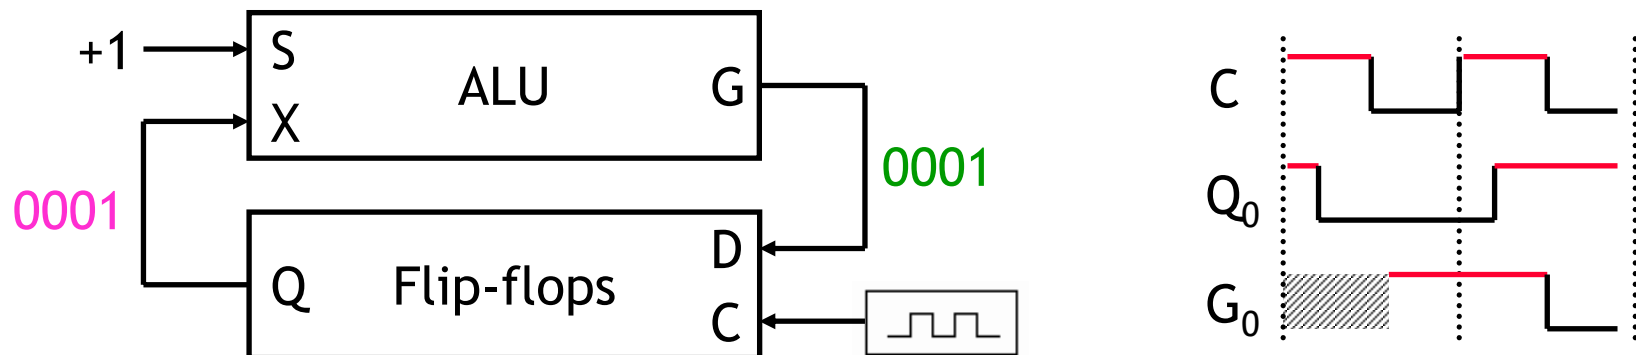


Example continued


- The ALU output is stored in the flip-flops on the next positive clock edge.



- The flip-flops “shut off” right after the positive clock edge, so no new values can be stored until the next positive edge.



Summary



- A **sequential circuit** has memory. It may respond differently to the same inputs, depending on its current **state**.
- **Latches** are the simplest memory units, storing individual bits. But it's difficult to control the timing of latches in a larger circuit.
- To use latches in bigger circuits, we need to do two things.
 - Keep the latches disabled until new values are ready to be stored.
 - Enable the latches just long enough for the update to occur.
- **Flip-flops** restrict memory writes to the positive edge of a **clock signal**.
- Next week we'll talk about how to analyze and design sequential circuits that use flip-flops as memory.