EEL 4783: HDL in Digital System Design

Lecture 5: Architeching Power

Prof. Mingjie Lin

Architecting Power

- The third of three primary physical characteristics of a digital design: power
- Relative to ASICs (application specific integrated circuits) with comparable functionality, FPGAs are power-hungry beasts and are typically not well suited for ultralow-power design techniques
- A number of FPGA vendors do offer low-power CPLDs (complex programmable logic devices), but these are very limited in size and capability and thus will not always fit an application that requires any respectable amount of computing power

Computing Power in CMOS

- In CMOS technology, dynamic power consumption is related to charging and discharging parasitic capacitances on gates and metal traces
- The general equation for current dissipation in a capacitor is I=V*C*f, where I is total current, V is voltage, C is capacitance, and f is frequency
- To reduce the current drawn, we must reduce one of the three key parameters
- In FPGA design, the voltage is usually fixed. This leaves the parameters C and f to manipulate the current. The capacitance C is directly related to the number of gates that are toggling at any given time and the lengths of the routes connecting the gates
- The frequency f is directly related to the clock frequency. All of the power-reduction techniques ultimately aim at reducing one of these two components.

Important Topics

- The impact of clock control on dynamic power consumption
- Problems with clock gating
- Managing clock skew on gated clocks
- Input control for power minimization
- Impact of the core voltage supply
- Guidelines for dual-edge triggered flip-flops
- Reducing static power dissipation in terminations

Clock Control

- The most effective and widely used technique for lowering the dynamic power dissipation in synchronous digital circuits is to dynamically disable the clock in specific regions that do not need to be active at particular stages in the data flow
- Most of the dynamic power consumption in an FPGA is directly related to the toggling of the system clock, temporarily stopping the clock in inactive regions of the design is the most straightforward method of minimizing this type of power consumption

How to Clock Gating?

- The recommended way to accomplish this is to use either the clock enable pin on the flip-flop or to use a global clock mux (in Xilinx devices this is the BUFGMUX element)
- If these clock control elements are not available in a particular technology, designers will sometimes resort to direct gating of the system clock
- Note that this is not recommended for FPGA designs, and this section describes the issues involved with direct gating of the system clock
- Clock control resources such as the clock enable flip-flop input or a global clock mux should be used in place of direct clock gating

Why is Clock Gating HARD?

- FPGAs are synchronous devices, and a number of difficulties arise when multiple domains are introduced through gating or asynchronous interfaces

- More in-depth discussion regarding clock domains later

Bad Example

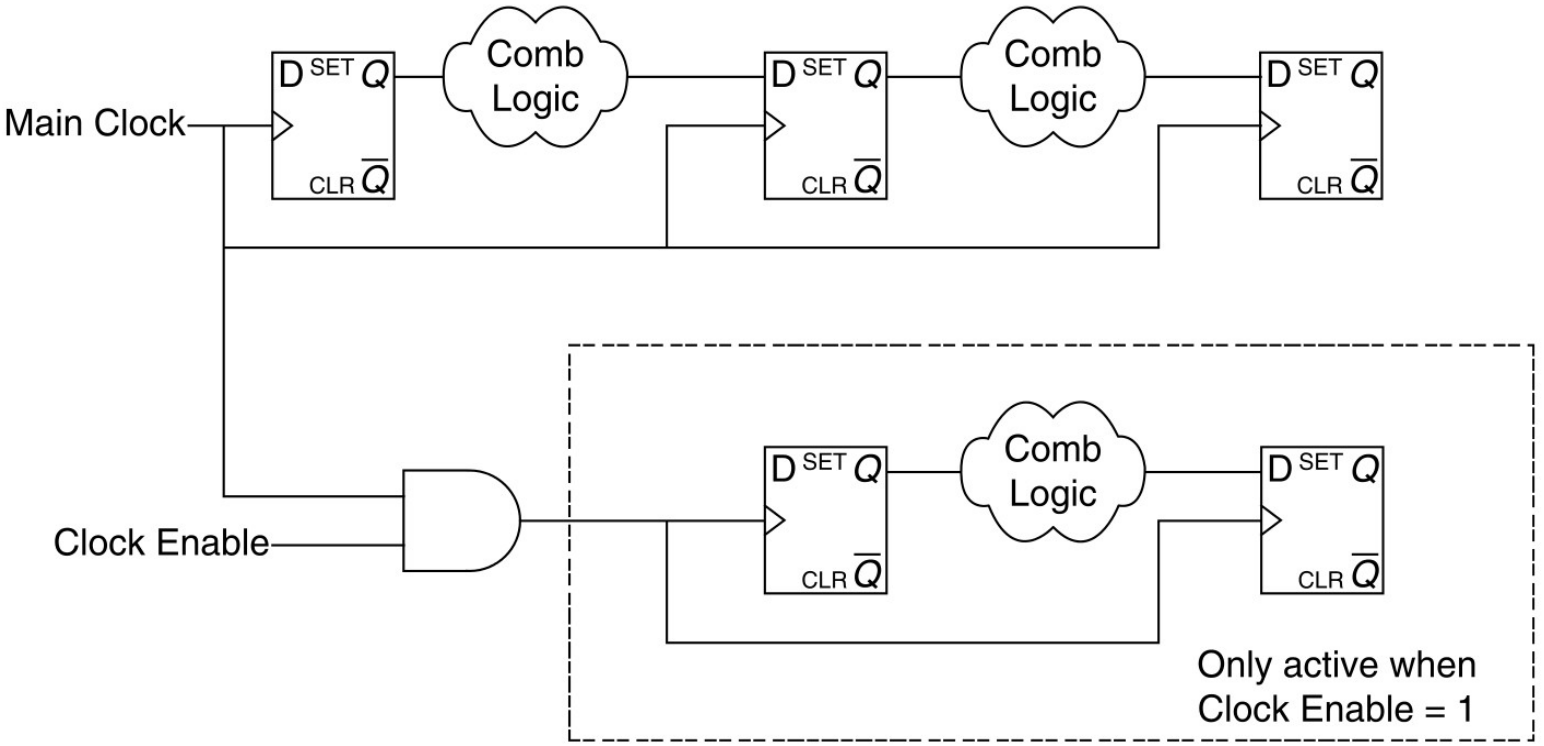

Clock gating is a direct means for reducing dynamic power dissipation but

creates difficulties in implementation and timing analysis

A gated clock introduces a new clock domain and will create difficulties for the FPGA designer

Clock Skew

- The concept of clock skew is a very important one in sequential logic design
- Mishandling clock skew can cause catastrophic failures in the FPGA

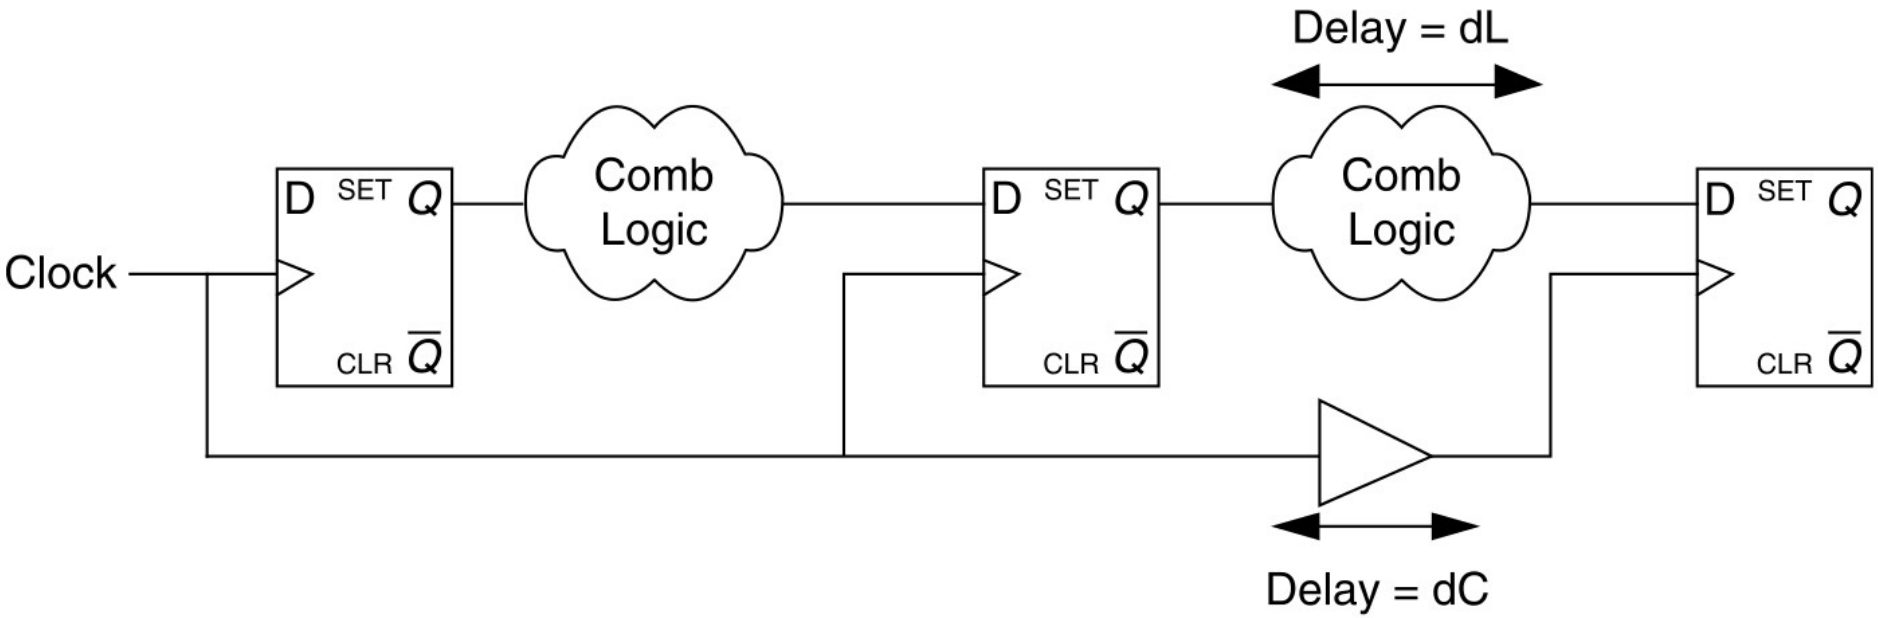

Managing Skew

- Low-skew resources provided on FPGAs ensure that the clock signal will be matched on all clock inputs as tightly as possible (within picoseconds)
- The clock line must be removed from the low-skew global resource and routed to the gating logic, in this case an AND gate

Final issues

- Please fill out the student info sheet before leaving
- Come by my office hours (right after class)
- Any questions or concerns?