

---

# Lecture 8:

## Clocking for High Performance Processors

Computer Systems Lab  
Stanford University  
horowitz@stanford.edu

Copyright © 2004 Mark Horowitz

---

EE371

Lecture 8-1

Horowitz

---

## Overview

### Introduction

In the last lecture we talked about different ways to build storage elements, and their characteristics. In this lecture we will look at how these storage elements interact with the design of logic and how it affects functionality and performance. Clocking and flop/latch design has a large influence on the circuit's power and performance. As was discussed in EE271, the role of the clocks is to keep signals correlated in time. There are many approaches to this problem, some even that don't require clocks. If clocks are used, it is critical to minimize the overhead caused by the clocks, which includes clock skew, and latch/flop overheads. This lecture looks at the sequencing issue and explores a couple of approaches to this problem.

---

EE371

Lecture 8-2

Horowitz

# Overview of Talk

---

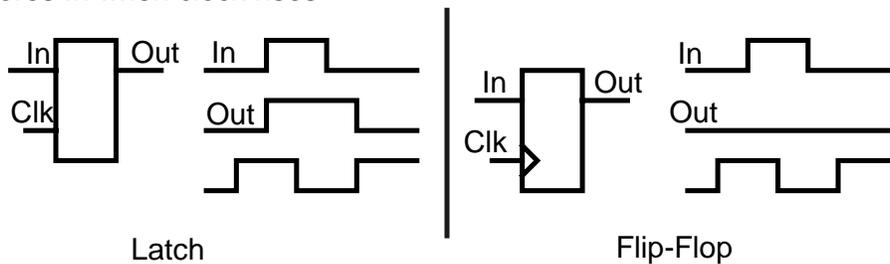
- Background:
  - Role of clocks
  - Why clocks are bad (clock overhead)
- Real world - (how to live with badness)
  - Clock distribution issues
  - Skew tolerant designs
- Summary

## Common View of Clock's Function

---

Clocks work with Latch or Flip-Flop to hold state

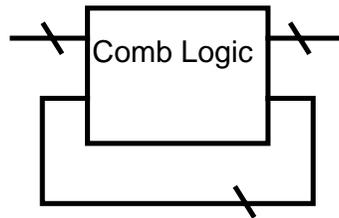
- Latch
  - Stores data when the clock is low
- Flip-Flop
  - Stores In when clock rises



## Another View

---

If the delay of every path **was EXACTLY** the same

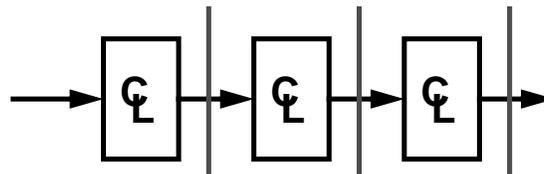


- I would not need clocks
  - The 'state' is stored in the gates and the wires.
- Signals stay naturally correlated in time
  - (wave pipelining)
- Impossible to do in practice, so ...

## Clock's Function:

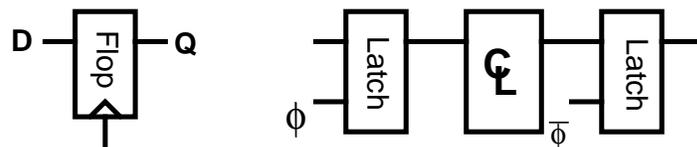
---

Keep values in a system correlated in time



Keep signals from racing ahead of others

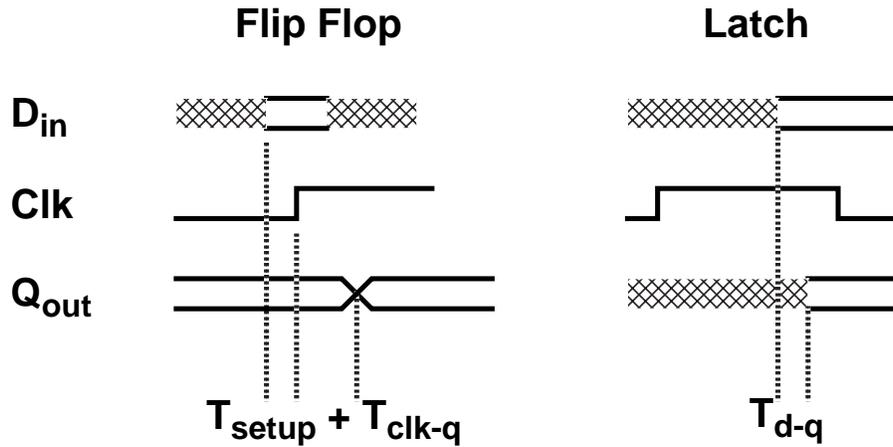
- Slow down signals that arrive too fast



- A flop is almost always built from two latches back to back

# Clock Overhead

Unfortunately, clocks delay slow paths too



# Clock Skew

Not all clocks arrive at the same time

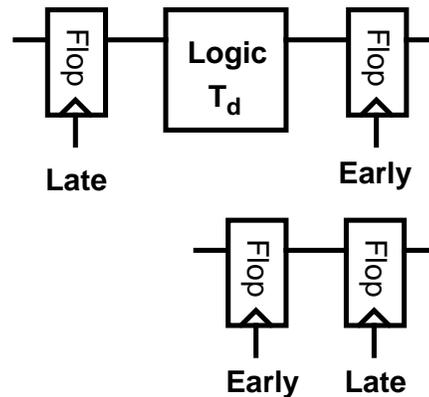
- Two problems:

- Adds more overhead:<sup>1</sup>

$$T_{cyc} = T_d + T_{setup} + T_{clk-q} + T_{skew}$$

- Can get the wrong answer:

$$T_{skew} < T_{clk-q} - T_{hold}$$



Low overhead -> Fast latches, low clock skew

1. As the previous lecture points out, it is hard to break the flop delay into a setup and clk-q delay, since the clk-q delay can increase when the setup time is small. For more precision, we need to consider the effective transparency window even for flops.

# Microprocessors

---

To understand why clock design is getting harder, all we need to do is look at a couple of processor designs over the past 20 years.

- Take a look at a few different processors
  - 8086 (1978)
  - R2000(1986)
  - 21064(1992)
  - 21164(1995)
  - 21264(1998)
  - P4(2001)
- Getting Larger
  - 30mm<sup>2</sup>, 80mm<sup>2</sup>, 220mm<sup>2</sup>, 300mm<sup>2</sup>, 300mm<sup>2</sup>, 220mm<sup>2</sup>
- Getting Faster
  - 5 MHz, 16MHz, 150MHz, 300MHz, 600MHz, 2GHz

# Processor Trends

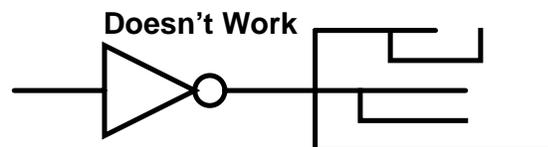
---

Performance:

- Part of speed increase is technology
  - Technology from 3 $\mu$  nMOS to 0.18 $\mu$  CMOS
  - CMOS FO4 gate delay is roughly 0.5L ns/ $\mu$
- Part is better circuit design
  - 200, 50, 20-25, 20, 16ish FO4 inv delays/cycle

Bottom Line:

- Cycle time getting shorter, even in # of gates
  - One FO4 delay is a larger % of cycle time
- Die size is growing
  - More capacitance on clocks (nF)
  - More resistance in clock lines



# Clock Overhead

---

Really two issues:

- Latch / flop delay
  - All latches I know of have a delay of  $\sim 1.5 FO4$
  - Two latches / cycle is  $> 15\%$  of cycle
- Clock skew
  - Keeping clock skew in ps constant is hard
  - But cycle times are falling
  - So engineering needed on the clock is growing.

If generating the clock is hard and getting harder, why do it

- Radical approach would be to eliminate it all together
  - Use local information instead
  - Called self-timed design
    - Uses information bundled with the signals for sequencing

# Self Timed Circuit Design Style

---

General idea:

- Instead of using (global) clock to sequence the data, use local information, such as the readiness of the input values and state of next stage logic, to provide sequencing of a pipelined logic

Benefits:

- No clock, or clock skew
- Circuits “runs at the speed it can run at”

New problems:

- In general, gathering the local information, making decisions based on that, and sending the decisions to the logic gates, take area, power and DELAY.
- Techniques to hide these “delay” uses delay matching techniques counting on matching.

Status quo will be here for a while

# Goal

---

**Be paranoid:**

**Make clock skew as small as possible**

**AND**

**Make your circuit insensitive to skew**

## Clock Distribution

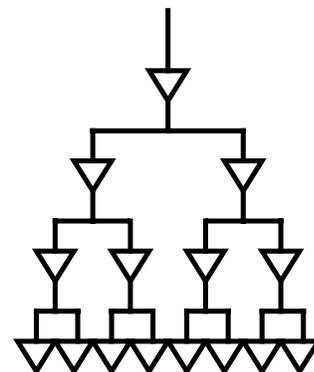
---

Need to reduce the skew on distributing the clock

- This requires us to reduce the wire delay, and the buffer delay
  - But we can't reduce the delay to the required levels (sub 100ps) so
- Make the effective delay small, by balancing the delays of all the paths
  - Change a total delay problem to a matching problem
  - Make  $\Delta T$  much smaller than  $T_{drive}$

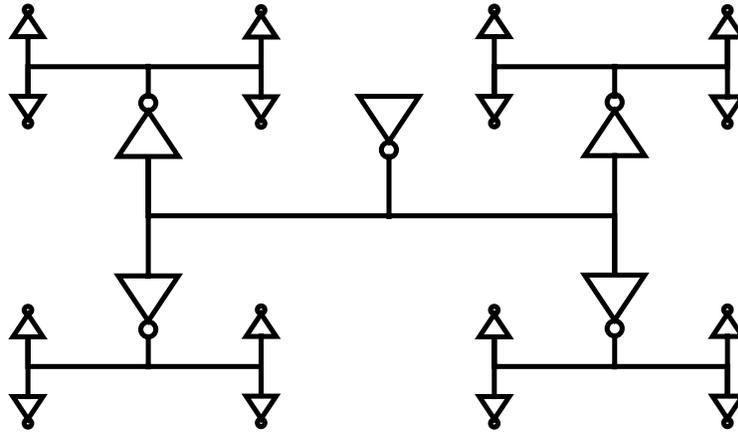
Use a clock trees

- Match the delay on different branches of tree
  - If the buffer delay matches
  - If the wire delay matches
  - Skew will be zero
- Obvious question:
  - How well can you match delays?



# H Trees

Space filling pattern that matches wire delays



Lots of papers on these things, but not real issue

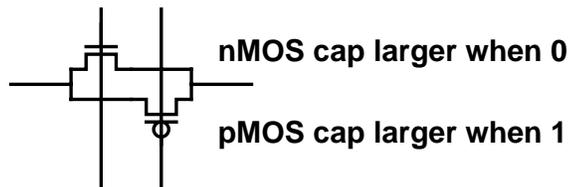
# Real Matching Issues

There are function blocks on chip, that mess up your nice abstract H-tree

- Wiring and buffers will need to fit
- Buffers eventually need to drive latches
  - Load of a latch is data dependent, since the source voltage can change

Gate loading depends on whether channel is formed

Variation depends on technology, but can be as much as 2:1 in capacitance



- Chip environments are not perfect
  - IR drops on the power supply lines
  - Temperature gradients across the chip
- Fabrication is not perfect
  - proximity effects / process tilt

# Wire Load Matching

---

Each wire has a different mix of components

- Not only gate-cap vs. wire cap
- Also % M1 - M2, M1 - M1, M1 fringe, etc.
- Need to find the worst-case skew

Process corners don't help

- Don't vary wires relative each other, don't account for data dependence

No real tools to help

- Problem for simulating matching of any kind, you need to simulate the worse-case for the matching, and this might not be either the slow or fast case.

Buffer Matching

- Buffer delay depend on Vdd, can vary over a chip due to IR drops (over 10%)
- Fabrication matching -- process tilt and proximity effects

# Process Tilt and Proximity Effects

---

Chips are large (2cm on a side), and transistors features are small

- Inverters on different sides on chip will be different
  - Difference is not in corner files, since corners make all transistors the same

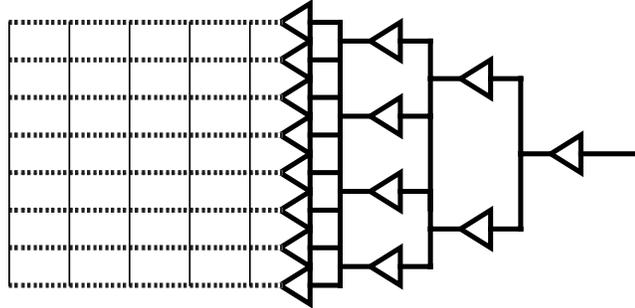
This data is not usually given to designers

It is essential for simulating clock skew

- Proximity Issues
  - Poly width sets channel length / speed of circuit
  - Current gate lengths are  $0.13\mu$ , Poly control must be within 10nm
  - Local poly environment affects etching rate, so it affects channel width
  - Matched inverters, need matched layout, and matched environments  
(region around buffer needs to be same -- add dummy buffers)

# Single Clock Distribution - 21064

- Thick metal layer for clocks, metal 3 ~2 $\mu$  thick
- Large clock buffer (entire vertical height of the chip)
  - Use a tree to balance the delay in this direction

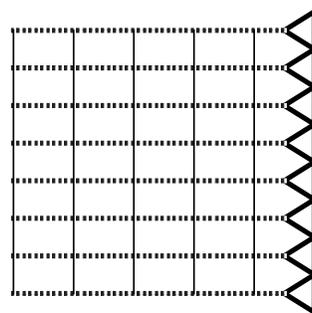
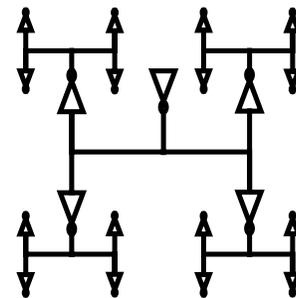


- Shorted together all the local clock wires
  - Main difference with a conventional tree; reduces the effects of mismatches
  - Especially effective for reducing local skew
- More recent processors have more clock buffers to keep skew small

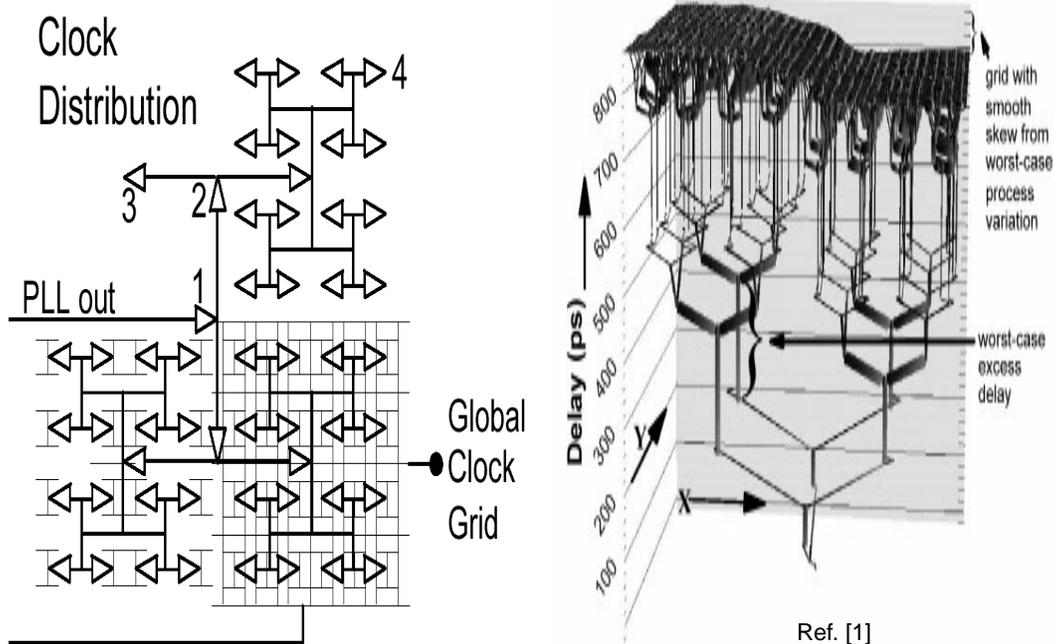
# Compare Trees and Grids

## Trees:

- Less routing resource
- Lower power
  - Less cap
  - Clock domain enable
- But, more “difficult design”
  - Uneven spacial distribution of load
  - “Last minute” load change
- Grids: (latches/flops loaded off the grid)
  - Lower skew
  - “Easier Design”
    - Regular routing environment
    - Earlier design constraints
  - But, routing resource and power implications



## Practical Design Example - Tree Grid Combo



EE371

Lecture 8-21

Horowitz

## Practical Design Example - Tree Grid Combo

Balanced tree to drive a grid for global distribution

- Minimizing global skew
- Easier design logistics
  - Pre-routes
  - Less sensitive to load variation
- Power “not too bad” since  $C_{\text{global}}$  is only a fraction of  $C_{\text{total}}$  for clocks

Local buffers, as loads for the global grid, can be conditioned and don't have to go long distance (hence easier skew control)

EE371

Lecture 8-22

Horowitz

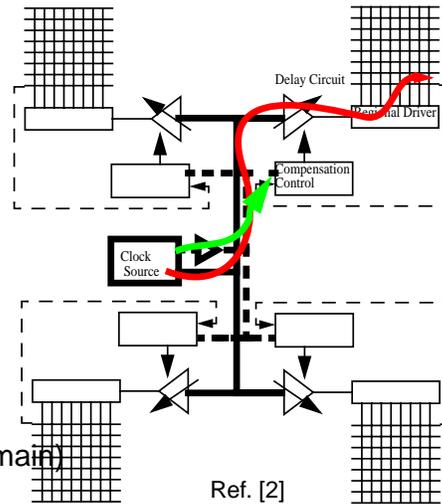
# Active Skew Compensation

Rational:

- Systematic in-die variation is accounted for
- Better productivity: Active compensation do not need details of lower level clock tree information

Residual error:

- Skew of ref. tree - Should be better
  - Shorter spacial span, fewer destinations
  - Symmetrical layout
  - No load variation
- Feedback clock error (skew within local domain)
  - Using a grid to minimize
- Phase detector error: Use smaller delay step, longer resolution time

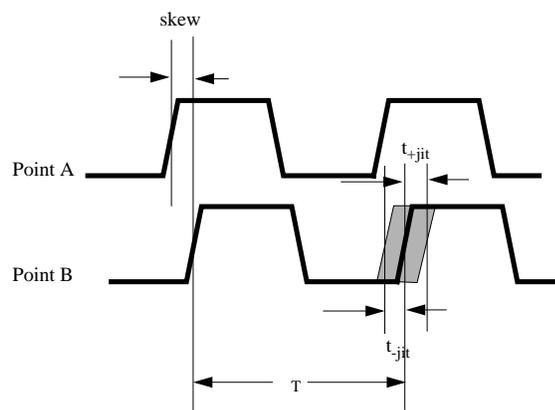


# Jitter in Clock Distribution

Jitter: Dynamic differences in the arrival time of the clock signals (same distribution point)

Short term jitter is of interest

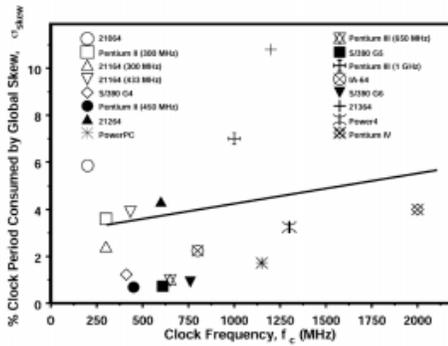
- Must be added to critical path budget



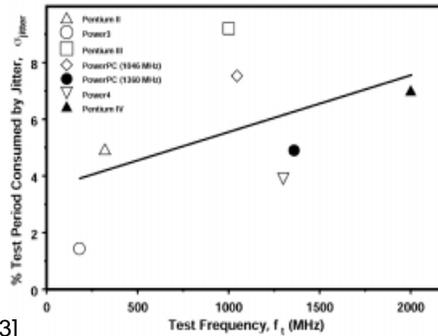
# Jitter in Clock Distribution

Two sources:

- From PLL and from clock distribution
  - PLL component, even as a percentage of cycle time, is relatively in check
  - Clock distribution component is going worse
  - Another incentive to minimize total delay

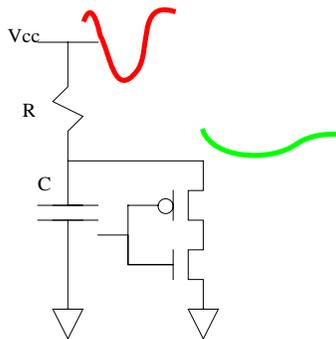


Ref. [3]



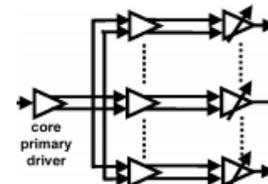
# Jitter Reduction

Power supply filters for clock buffers



Differential clocking

- Less sensitive to ground/vdd bounces
- Pay in area/effort, so only used close to the root of the tree



# Physical Design for Minimal Skew and Jitter

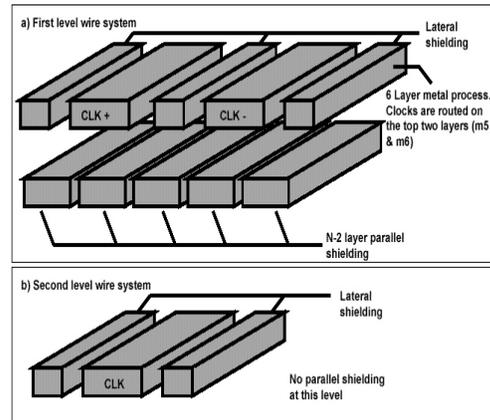
Matching (everything)

Active circuits:

- Same cell template
- Placement orientation
- “Environment” (Proximity effect)

Wires:

- Full shields
  - Less capacitive coupling
  - Inductive coupling important too
    - **Line impedance matching starts to be a factor**



Ref. [4]

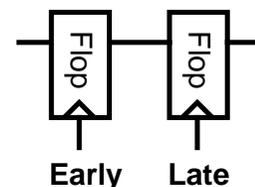
## Local Skew

Important for race-through

- Can get the wrong answer:
  - $T_{skew} < T_{clk-q} - T_{hold}$

Only occurs if delay is less than skew

- Delay is small only when elements are close
  - So only occurs when local skew is large
  - Shorting clocks together can reduce local skew
    - But also limits your design
    - Only have one clock and it is on all the time
    - Not the lowest power solution
    - Can avoid this problem by not having short paths
  - Need to have tool to check (and fix) min-delay
    - Delay cell insertion



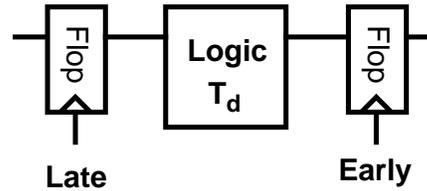
## Global Skew

---

Still have long path problem

- Skew adds more overhead:

- $T_{cyc} = T_d + T_{setup} + T_{clk-q} + T_{skew}$



So don't use flops!

- The situation with latches is a little different: Transparency windows
- Clock skew does not have to impact  $T_{cyc}$  if data arrives in the middle of a transparency window

## Clocking Design

---

- Trade off between overhead / robustness / complexity

Constraints on the logic

vs.

Constraints on the clocks

- For performance, you need to worry about

- Overhead of the sequencing

Delay through the latches/flops

Wasted time from clock-skew

- Look at a number of different clocking methods:

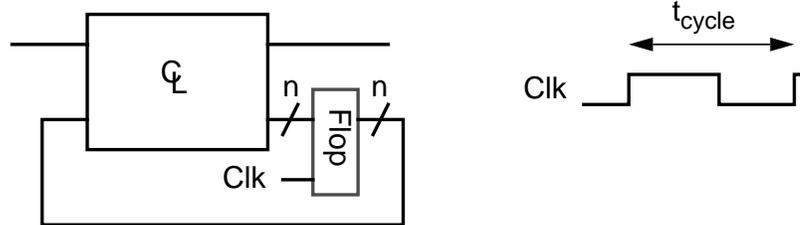
Edge triggered clocking

Pulse mode clocking

Two phase clocking (might only have one clock)

# Edge Triggered Flop Design

- Most popular design style (comes from old TTL designs)
- Used in many ASIC designs (Gate Arrays and standard Cells)
- Using a single clock, breaks every cycle with a flip-flop



- Timing Constraints

$$t_{\text{dmax}} < t_{\text{cycle}} - t_{\text{setup}} - t_{\text{clk-q}} - t_{\text{skew}}$$

$$t_{\text{dmin}} > t_{\text{skew}} + t_{\text{hold}} - t_{\text{clk-q}}$$

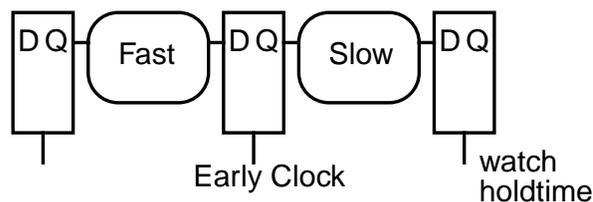
- If skew is large enough, you have two sided timing constraints

# Flop Design

Flops introduce hard timing boundaries in to the circuit

- Data must setup before the clock edge
- Output does not change until after the clock edge
  - Any uncertainty in clock, or data is wasted
  - Need to know precisely when the data will arrive

If some section of logic will be done early, to use that extra time, you need to move the clock to the flop early, so the next cycle has more time (and you had better check the hold time of the following flop)



# Practical Issues in Skew Analysis/Clock Tweaking

---

In reality, moving clock at design phase for long path is done but not too often

- Short pathological path from/to the same clock exist
- There are modeling/process errors for delay

Most common trick: design clock buffers whose delay can be tuned on silicon

- Commit only what are verified to be helpful for the majority of silicon
- Area overhead; May not have the right granularity

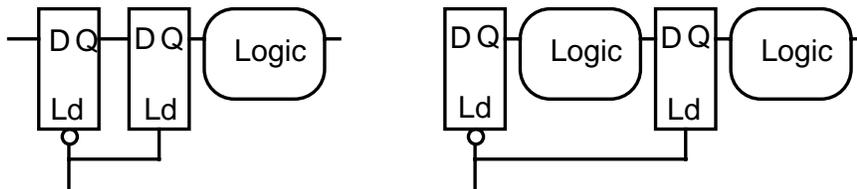
Given the modeling error for skews and the difficulty to get those numbers during the design phase, in doing timing analysis, absolute skew numbers are typically not used, a systematic skew budgeting is typically used. An example:

- Path between sequential element driven by clocks from the same local buffer
- Path between sequential element driven by clocks from the different local buffer but same regional buffer
- Path between sequential element driven by clocks from the different regional buffer

## Latch Based Design

---

- Break flop into its two latches, and place logic between the latches.



- There are no hard boundaries in latches
  - Pass data when clock is high
- Latching event is the load to hold transition
  - If data arrives early it is passed through
  - Can borrow time naturally, and
  - Is insensitive to clock skew, for critical paths, data sets the timing (well generally)

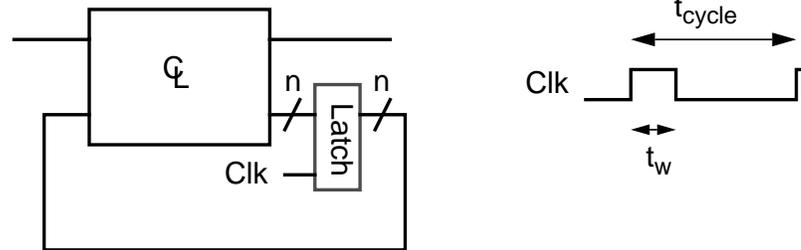
## Pulse Mode Clocking

---

Two requirements:

- All loops of logic are broken by a single latch
- The clock is a narrow pulse

It must be shorter than the shortest path through the logic



Clock is usually generated inside the latch

- Timing Requirements

$$t_{dmax} < t_{cycle} - t_{d-q} - t_{skew}$$

$$t_{dmin} > t_w - t_{d-q} + t_{skew}$$

## Pulse Mode Clocking

---

- Used in the original Cray computers (ECL machines)
- Advantage is it has a very small clocking overhead
  - One latch delay added to cycle
- Leads to double sided timing constraints
  - If logic is too slow OR too fast, the system will fail
  - But there is some flow time when the latch is enabled (softer edge)
- Pulse width is critical
  - Hard to maintain narrow pulses through inverter chains
- People use this type of clocking for high speed MOS circuits
  - Pulse generation is done in each latch.
  - Clock distributed is 50% duty cycle
  - CAD tools check min delay
  - Called a glitch flop, but it is not a flop, it is a glitch latch!

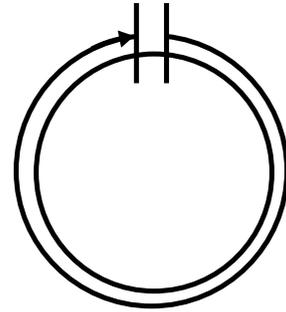
# Thinking About Timing

Imagine your arranging your netlist on a sheet where all the flops at the top

- The gates distance from the top indicates the settling time of its output
- Gates at the end of long paths would be at the bottom of the sheet
- Some of the outputs are the inputs to the flops, so we roll the sheet
- Forms a cylinder, where the circumference is equal to the cycle time

With flops, the input has to settle before its clock rises, and the output can't change until its clock falls so

- To guarantee operation, need to waste skew time
- Hard edge is the problem



# Latches

Are hard to analyze, since the timing of the output is not completely set by clock

- Output is valid  $clk-q$  delay after clock rises  
If input was valid when clock was low
- Output is valid a  $d-q$  delay after input  
If input becomes valid when clock is high

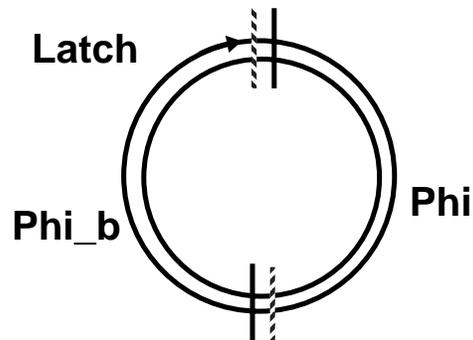
Skew changes the relative timing of clock and data

- Since latches are soft barriers it does not change the output arrival time!
- Latch based system can tolerate skew

$$= T_{cyc} - T_{max}$$

$$T_{cyc} = \text{cycle time}$$

$$T_{max} = \text{Max delay between latches}$$



# Summary

---

Clocks have skew

- Today skew is caused by mismatches of balanced paths
  - Simulating matching is VERY HARD since need data that is not available
  - Need your our simulation environment since assuming perfect matching gives you get best-case skew
- I believe that keeping global skew under 100ps is hard
  - Can do better than that over smaller regions (for fixed functionality, skew should roughly scale with inverter speed)
- Need to have circuits that deal with skew
  - Need to prevent race-through by padding short paths (and min. local skew)
  - Prevent cycle-time impact by using latch based design techniques
    - Use skew tolerant design technique

If you are careful clocks will work fine for chips running at GHz rates.

## References:

### [1] The clock distribution of the Power4 microprocessor

Restle, P.J.; Carter, C.A.; Eckhardt, J.P.; Krauter, B.L.; McCredie, B.D.; Jenkins, K.A.; Weger, A.J.; Mule, A.V.; Solid-State Circuits Conference, 2002. Digest of Technical Papers. ISSCC. 2002 IEEE International, Volume: 1, 3-7 Feb. 2002 Pages:144 - 145 vol.1

### [2] Clock generation and distribution for the first IA-64 microprocessor

Tam, S.; Rusu, S.; Nagarji Desai, U.; Kim, R.; Ji Zhang; Young, I.; Solid-State Circuits, IEEE Journal of, Volume: 35, Issue: 11, Nov. 2000 Pages:1545 - 1552

[3] Electrical and optical clock distribution networks for gigascale microprocessors Mule, A.V.; Glytys, E.N.; Gaylord, T.K.; Meindl, J.D.; Very Large Scale Integration (VLSI) Systems, IEEE Transactions on, Volume: 10, Issue: 5, Oct. 2002 Pages:582 - 594

[4] The core clock system on the next generation Itanium1 microprocessor Anderson, F.E.; Wells, J.S.; Berta, E.Z.; Solid-State Circuits Conference, 2002. Digest of Technical Papers. ISSCC. 2002 IEEE International, Volume: 1, 3-7 Feb. 2002 Pages:146 - 453 vol.1