



EECS151/251A

Spring 2021

Digital Design and Integrated Circuits

Instructor:
John Wawrzynek

Lecture 24:

Clocks, Packaging, and Power
Distribution

Announcements

- ❑ Virtual Front Row for today 4/20:
 - ❑ *Jose Rodriguez*
 - ❑ *Rahul Arya*
 - ❑ *Zitao Fang*
 - ❑ *Jeremy Ferguson*
- ❑ **Please ask question or make comments!**
- ❑ Homework assignment 10 posted - due next Monday.
- ❑ HW 11 - final problem set - posted end of this week.

Announcements

□ End game:

14	4/20	Clock and Power Distribution		Checkpoint 2 due, , Checkpoint 3 & 4 Released	
	4/22	Testing, Faults, Error Correction Codes			
15	4/27	Inside Logic Synthesis Tools		Checkpoint 3 due	
	4/29	Wrap-up and Exam Review			
16	5/5	RRR No Lecture		Checkpoint 4	Final Checkoff due
	5/7	RRR No Lecture			
FINAL	5/14	No Class - Final Exam 7-10 PM			

Outline




- ❑ Clock non-idealities
- ❑ Clock Distribution
- ❑ Chip packaging
- ❑ Power Distribution




Synchronous Timing - Review

Synchronous Timing

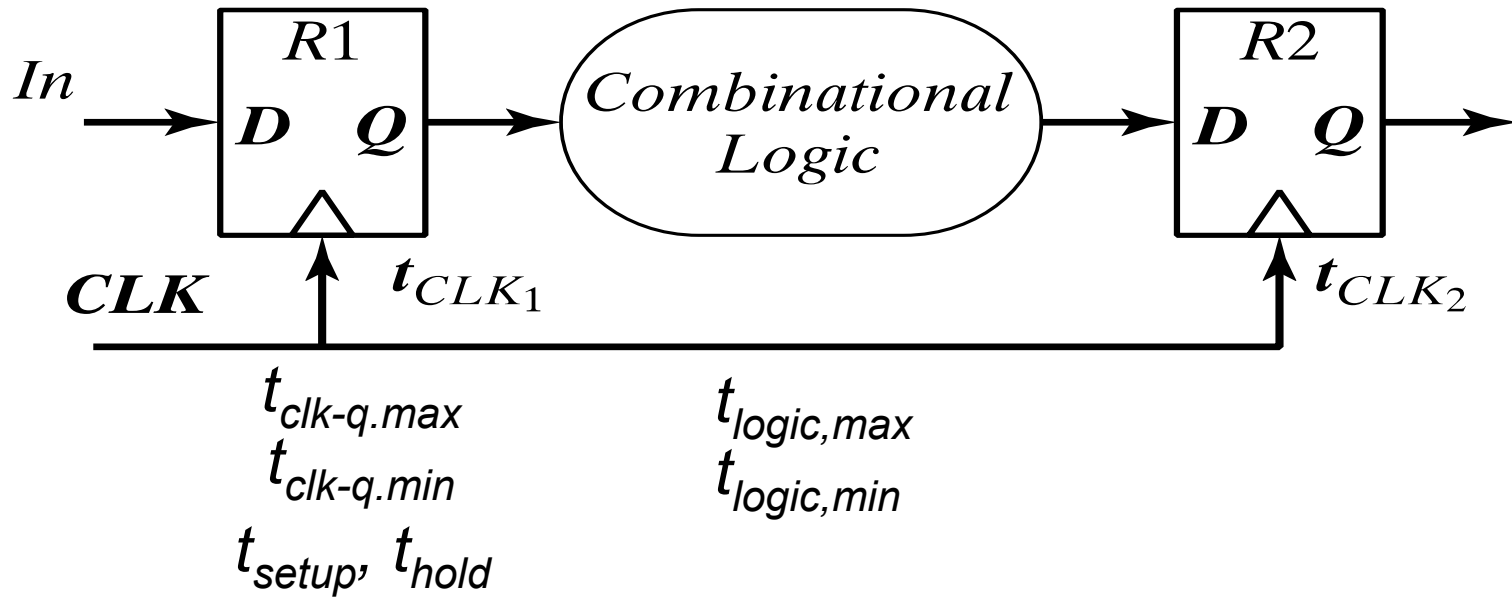


Register Timing Parameters

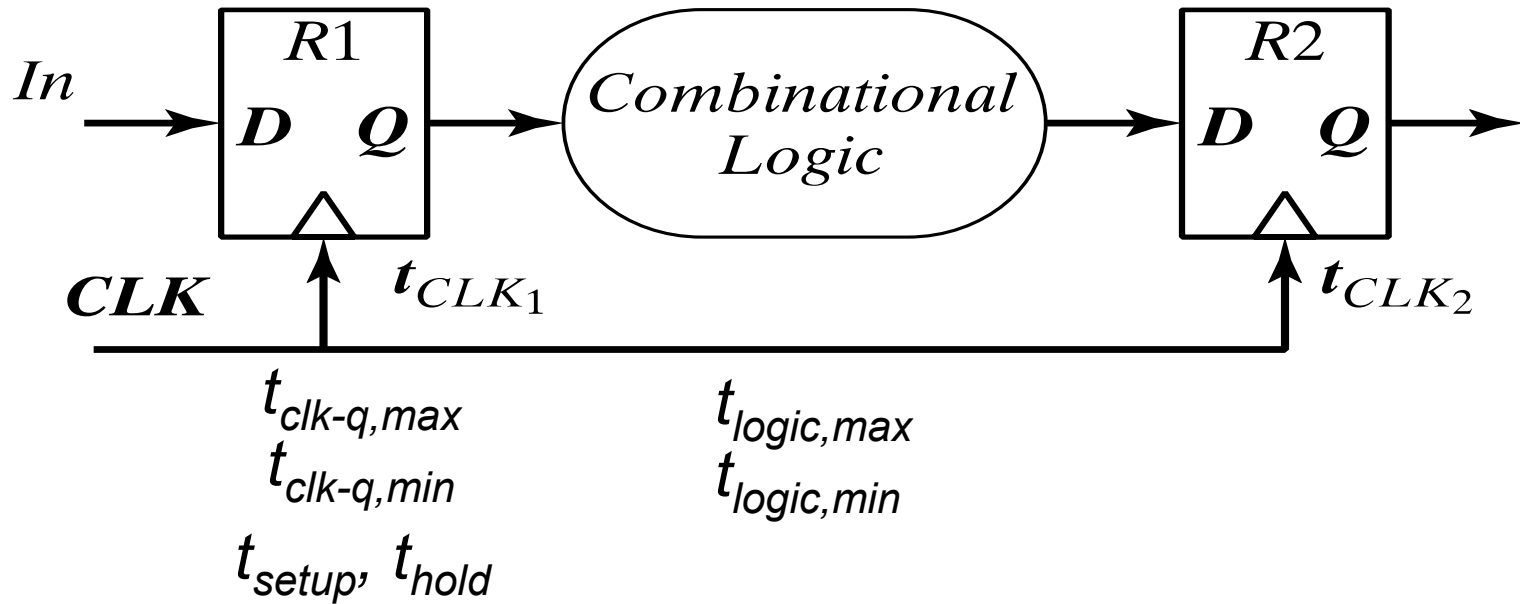


Output delays can be different for rising and falling data transitions

Timing Constraints



Timing Constraints



Cycle time (max): $T_{clk} > t_{clk-q,max} + t_{logic,max} + t_{setup}$

Race margin (min): $t_{hold} < t_{clk-q,min} + t_{logic,min}$



Clock Nonidealities

Clock Nonidealities

❑ Clock skew: t_{SK}


- Time difference between the sink (receiving) and source (launching) clock edge; deterministic + random

❑ Clock jitter

- Temporal variations in consecutive edges of the clock signal; modulation + random noise
- Cycle-to-cycle (short-term) t_{JS}
- Long term t_{JL}

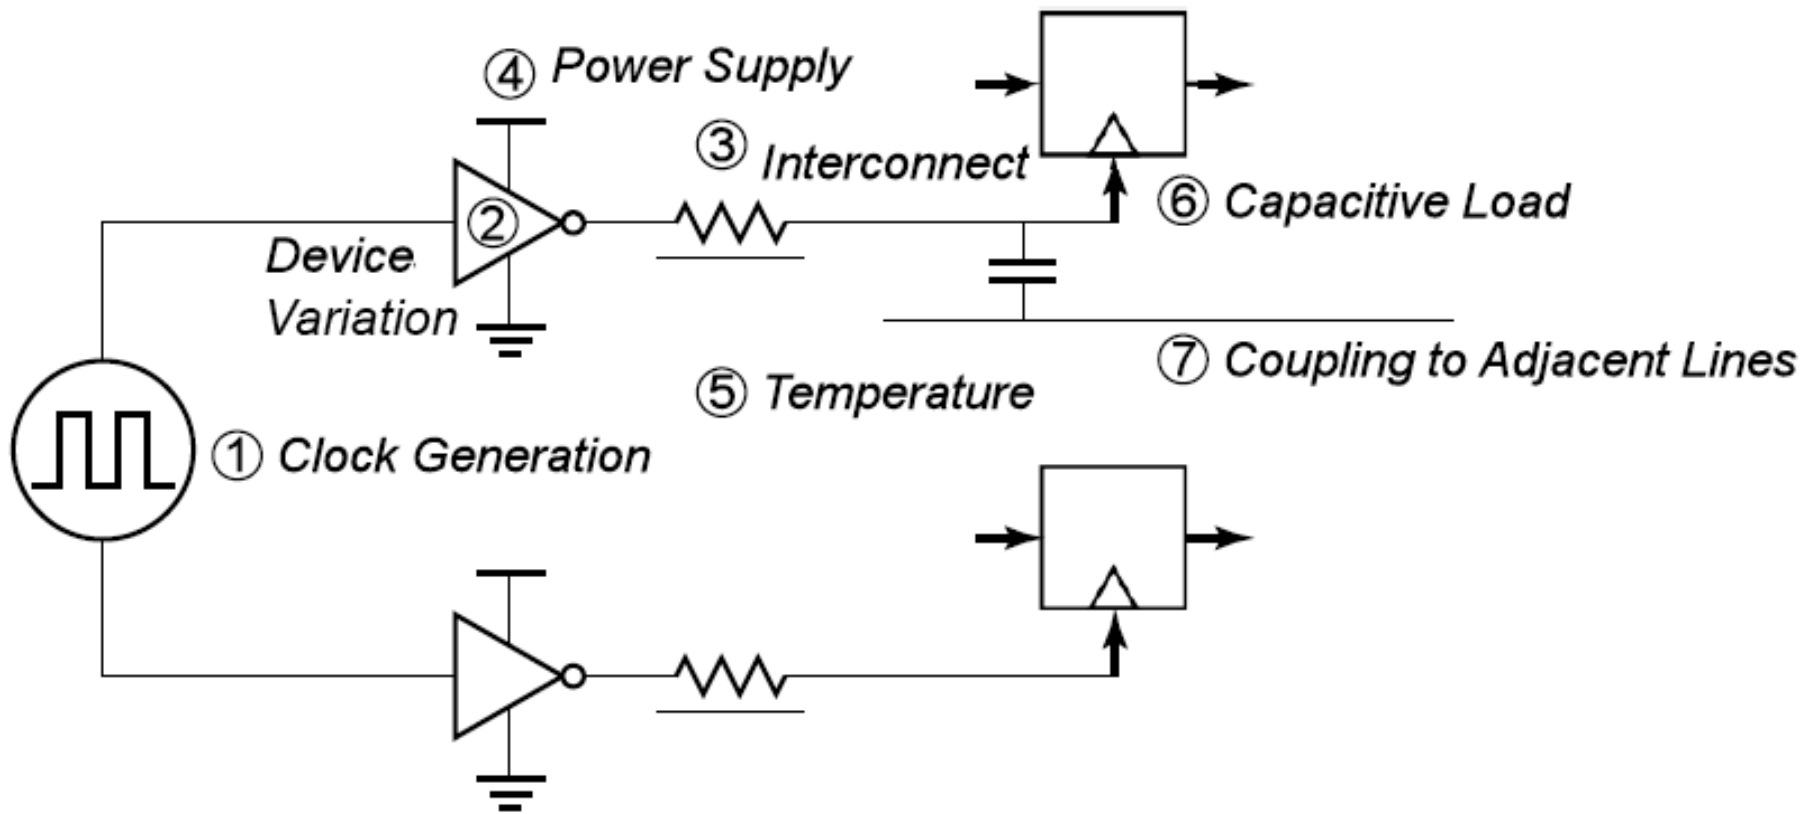
❑ Variation of the pulse width

- Important for level sensitive clocking




Clock Uncertainties

Sources of clock uncertainty

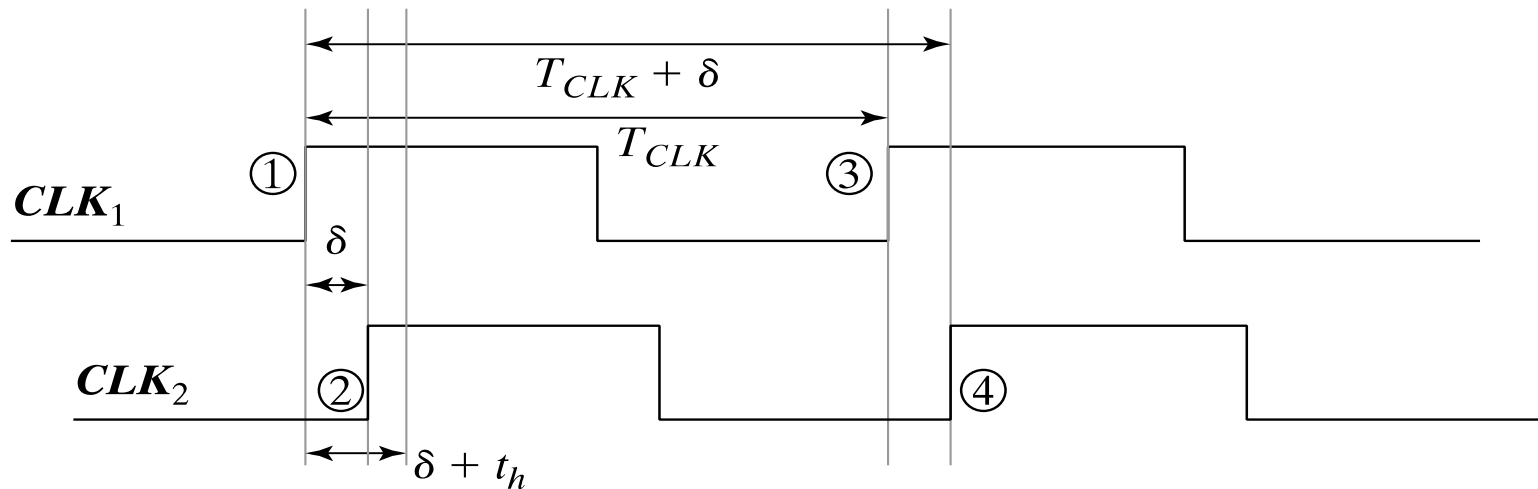



Clock Skew and Jitter



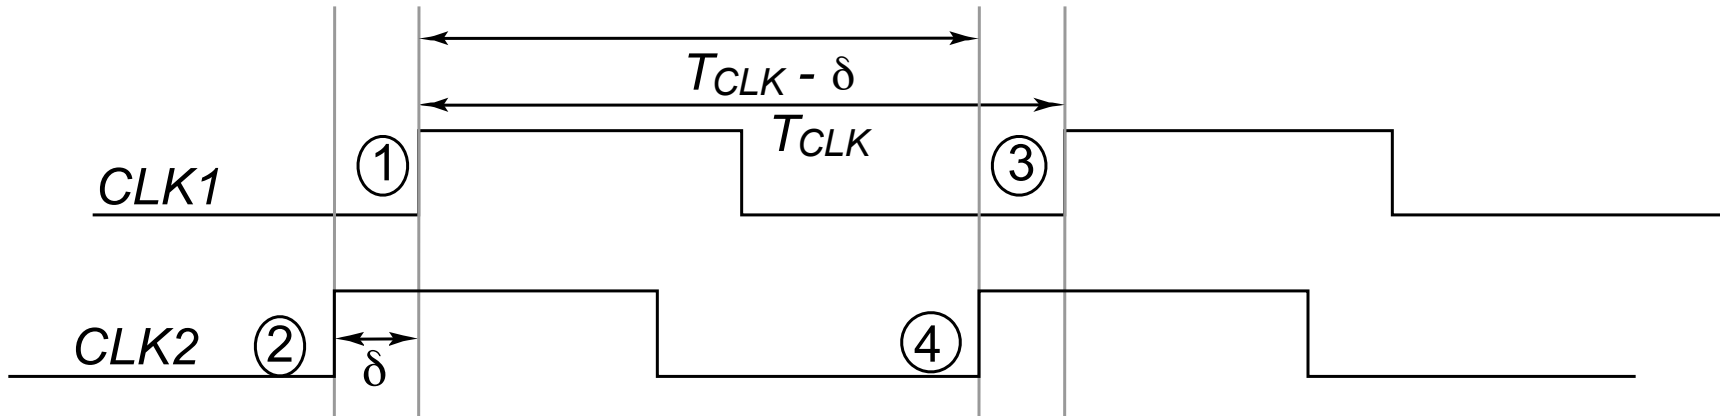
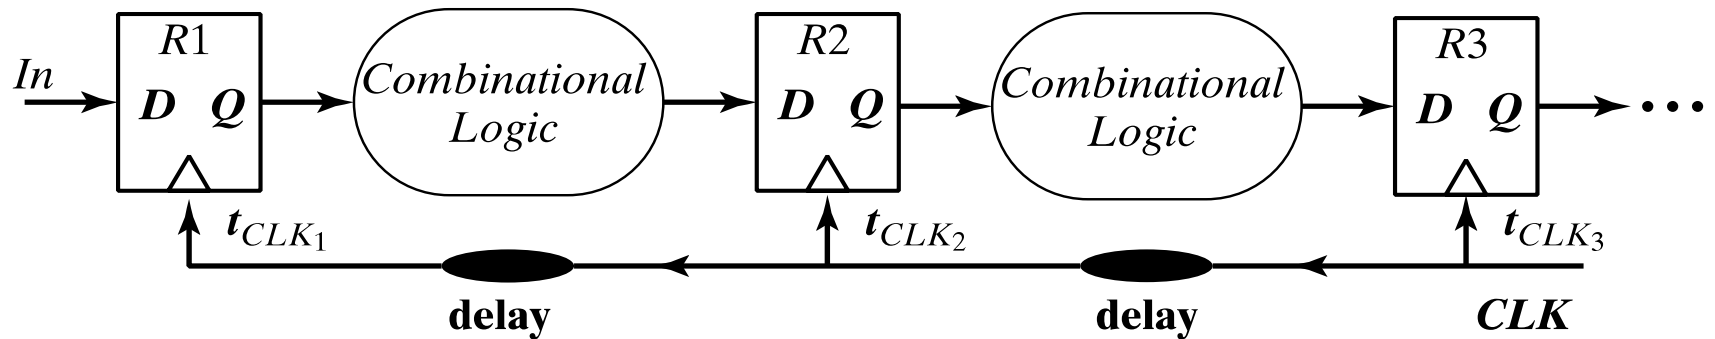
- Both skew and jitter affect the effective cycle time and the race margin

Positive Skew




Launching edge arrives before the receiving edge

Negative Skew



Receiving edge arrives before the launching edge

Timing Constraints



$t_{clk-q,max}$
 $t_{clk-q,min}$
 t_{setup}, t_{hold}


$t_{logic,max}$
 $t_{logic,min}$

Minimum cycle time:

$$T_{clk} + \delta = t_{clk-q,max} + t_{setup} + t_{logic,max}$$

Skew may be negative or positive

Timing Constraints



$t_{clk-q,max}$
 $t_{clk-q,min}$
 t_{setup}, t_{hold}

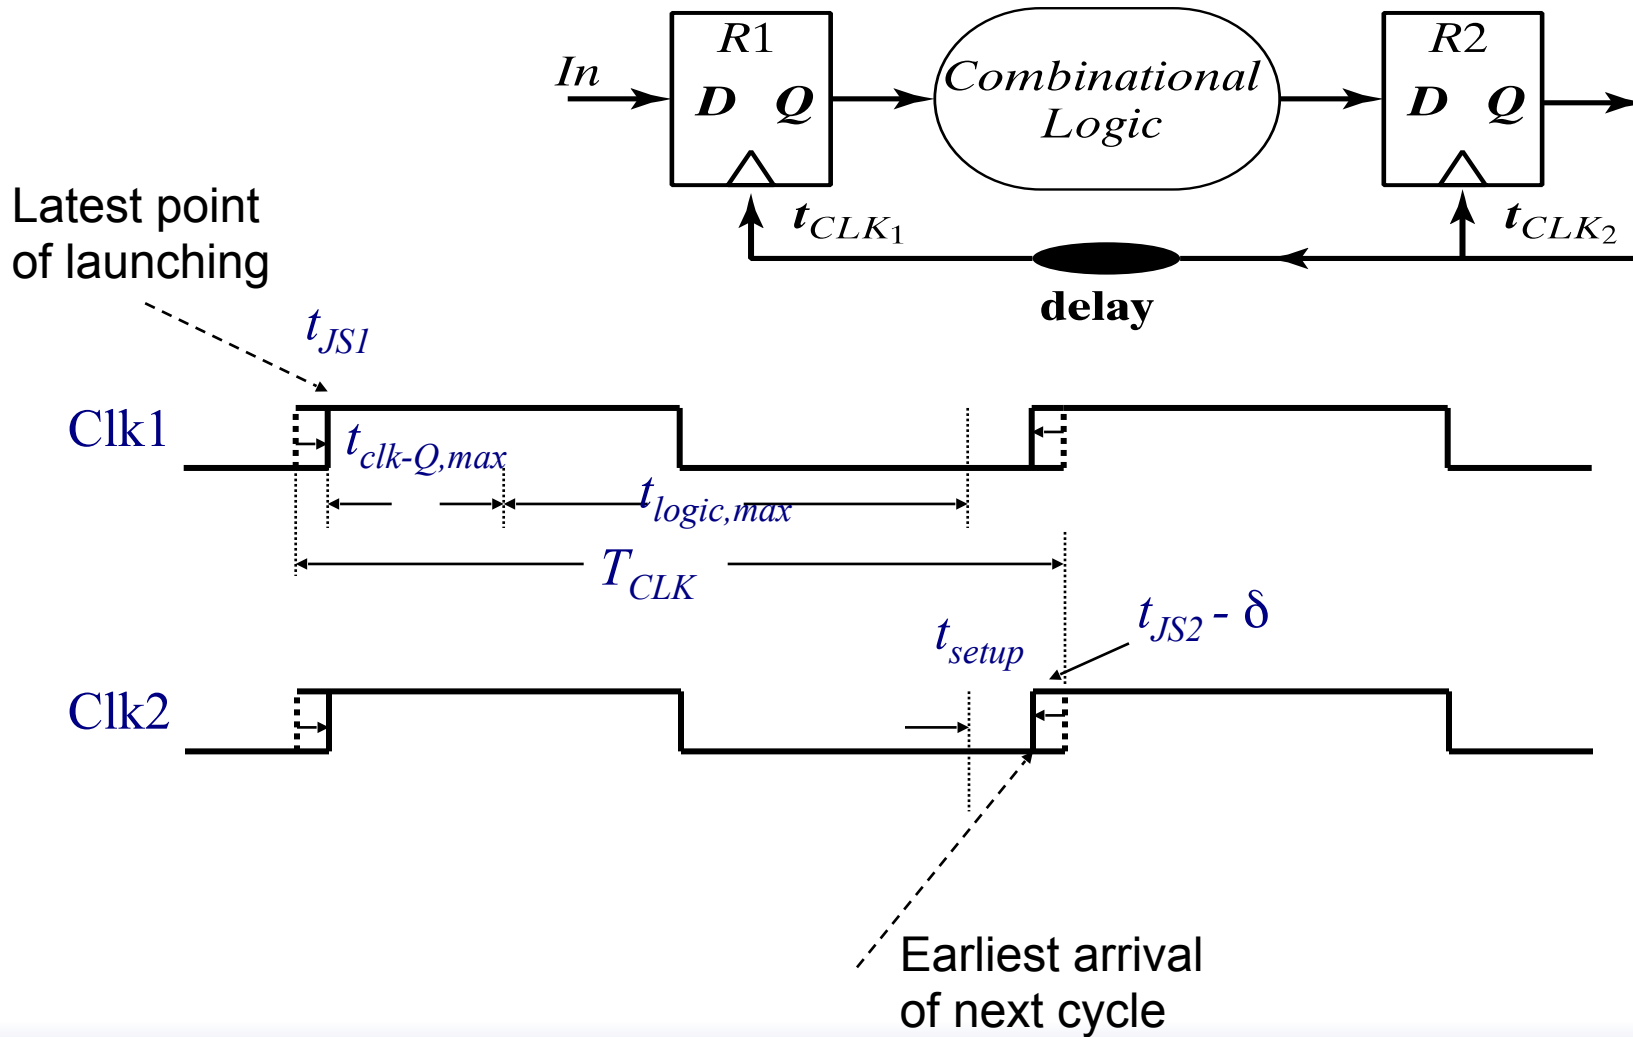
$t_{logic,max}$
 $t_{logic,min}$

Hold time constraint:

$$t_{(clk-q,min)} + t_{(logic,min)} > t_{hold} + \delta$$

Skew may be negative or positive

Jitter Contributes to Critical Path



Clock Constraints in Edge-Triggered Systems

If launching edge is late and receiving edge is early, the data will not be too late if:

$$t_{clk-q,max} + t_{logic,max} + t_{setup} < T_{CLK} - t_{JS,1} - t_{JS,2} + \delta$$

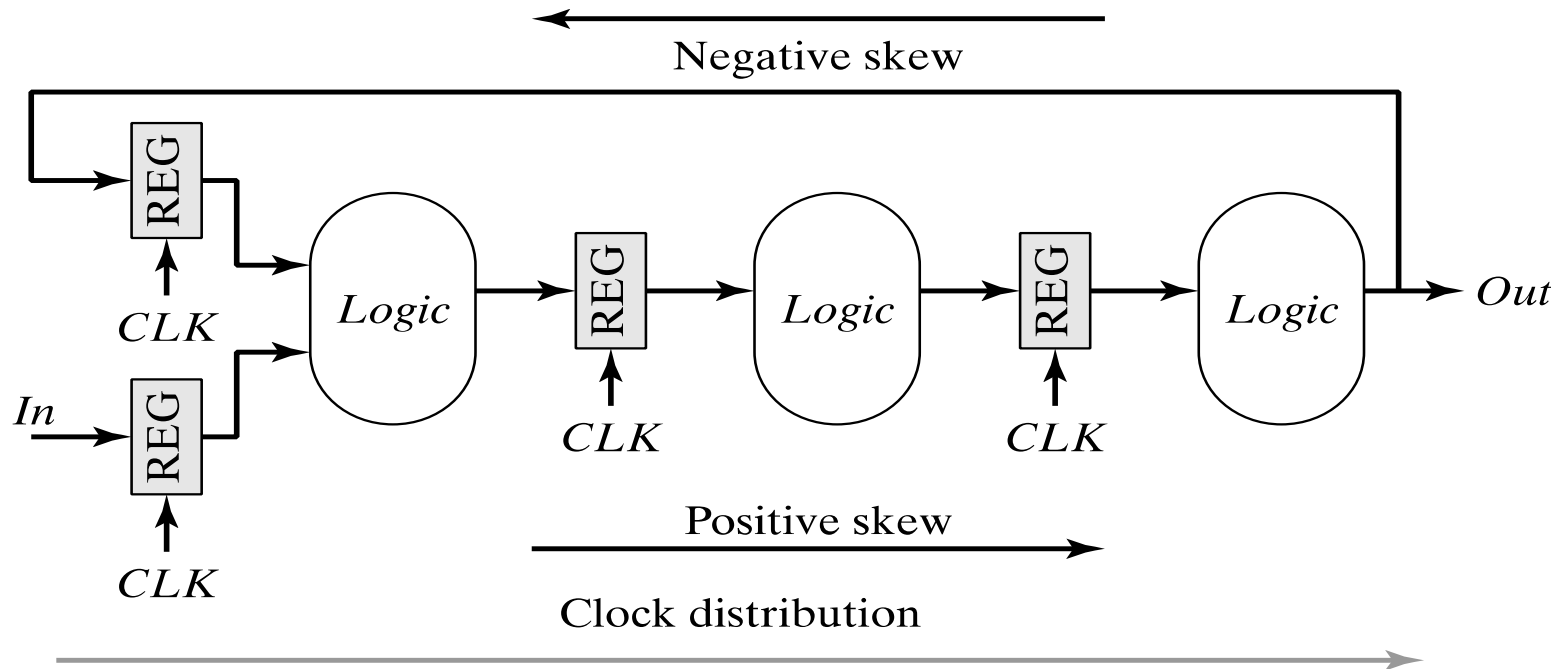
Minimum cycle time is determined by the maximum delays through the logic

$$t_{clk-q,max} + t_{logic,max} + t_{setup} - \delta + 2t_{JS} < T_{CLK}$$

Skew can be either positive or negative

Jitter t_{JS} usually expressed as peak-to-peak or $n \times$ RMS value

Datapath with Feedback





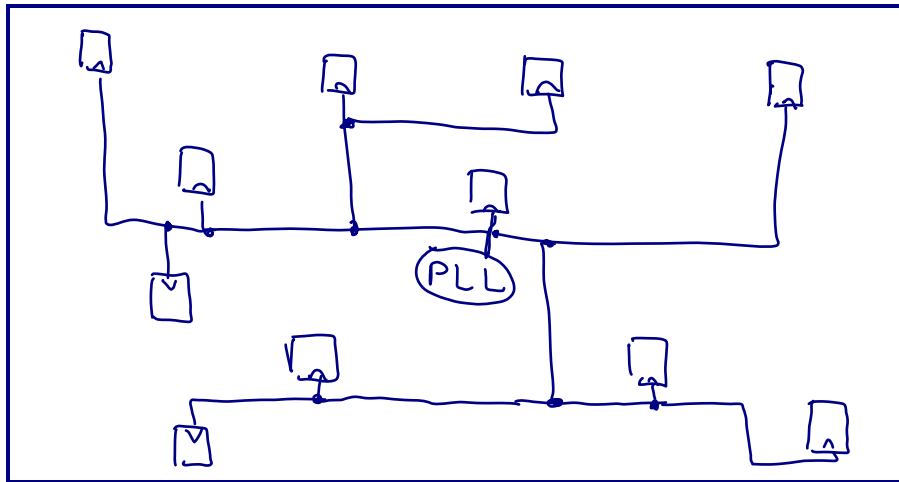
Clock Distribution

Clock Distribution



- Single clock generally used to synchronize all logic on the same chip (or region of chip)
 - Need to distribute clock over the entire region
 - While maintaining low skew/jitter
 - And without burning too much power

Clock Distribution

- ❑ What's wrong with just routing wires to every point that needs a clock?




H-Tree

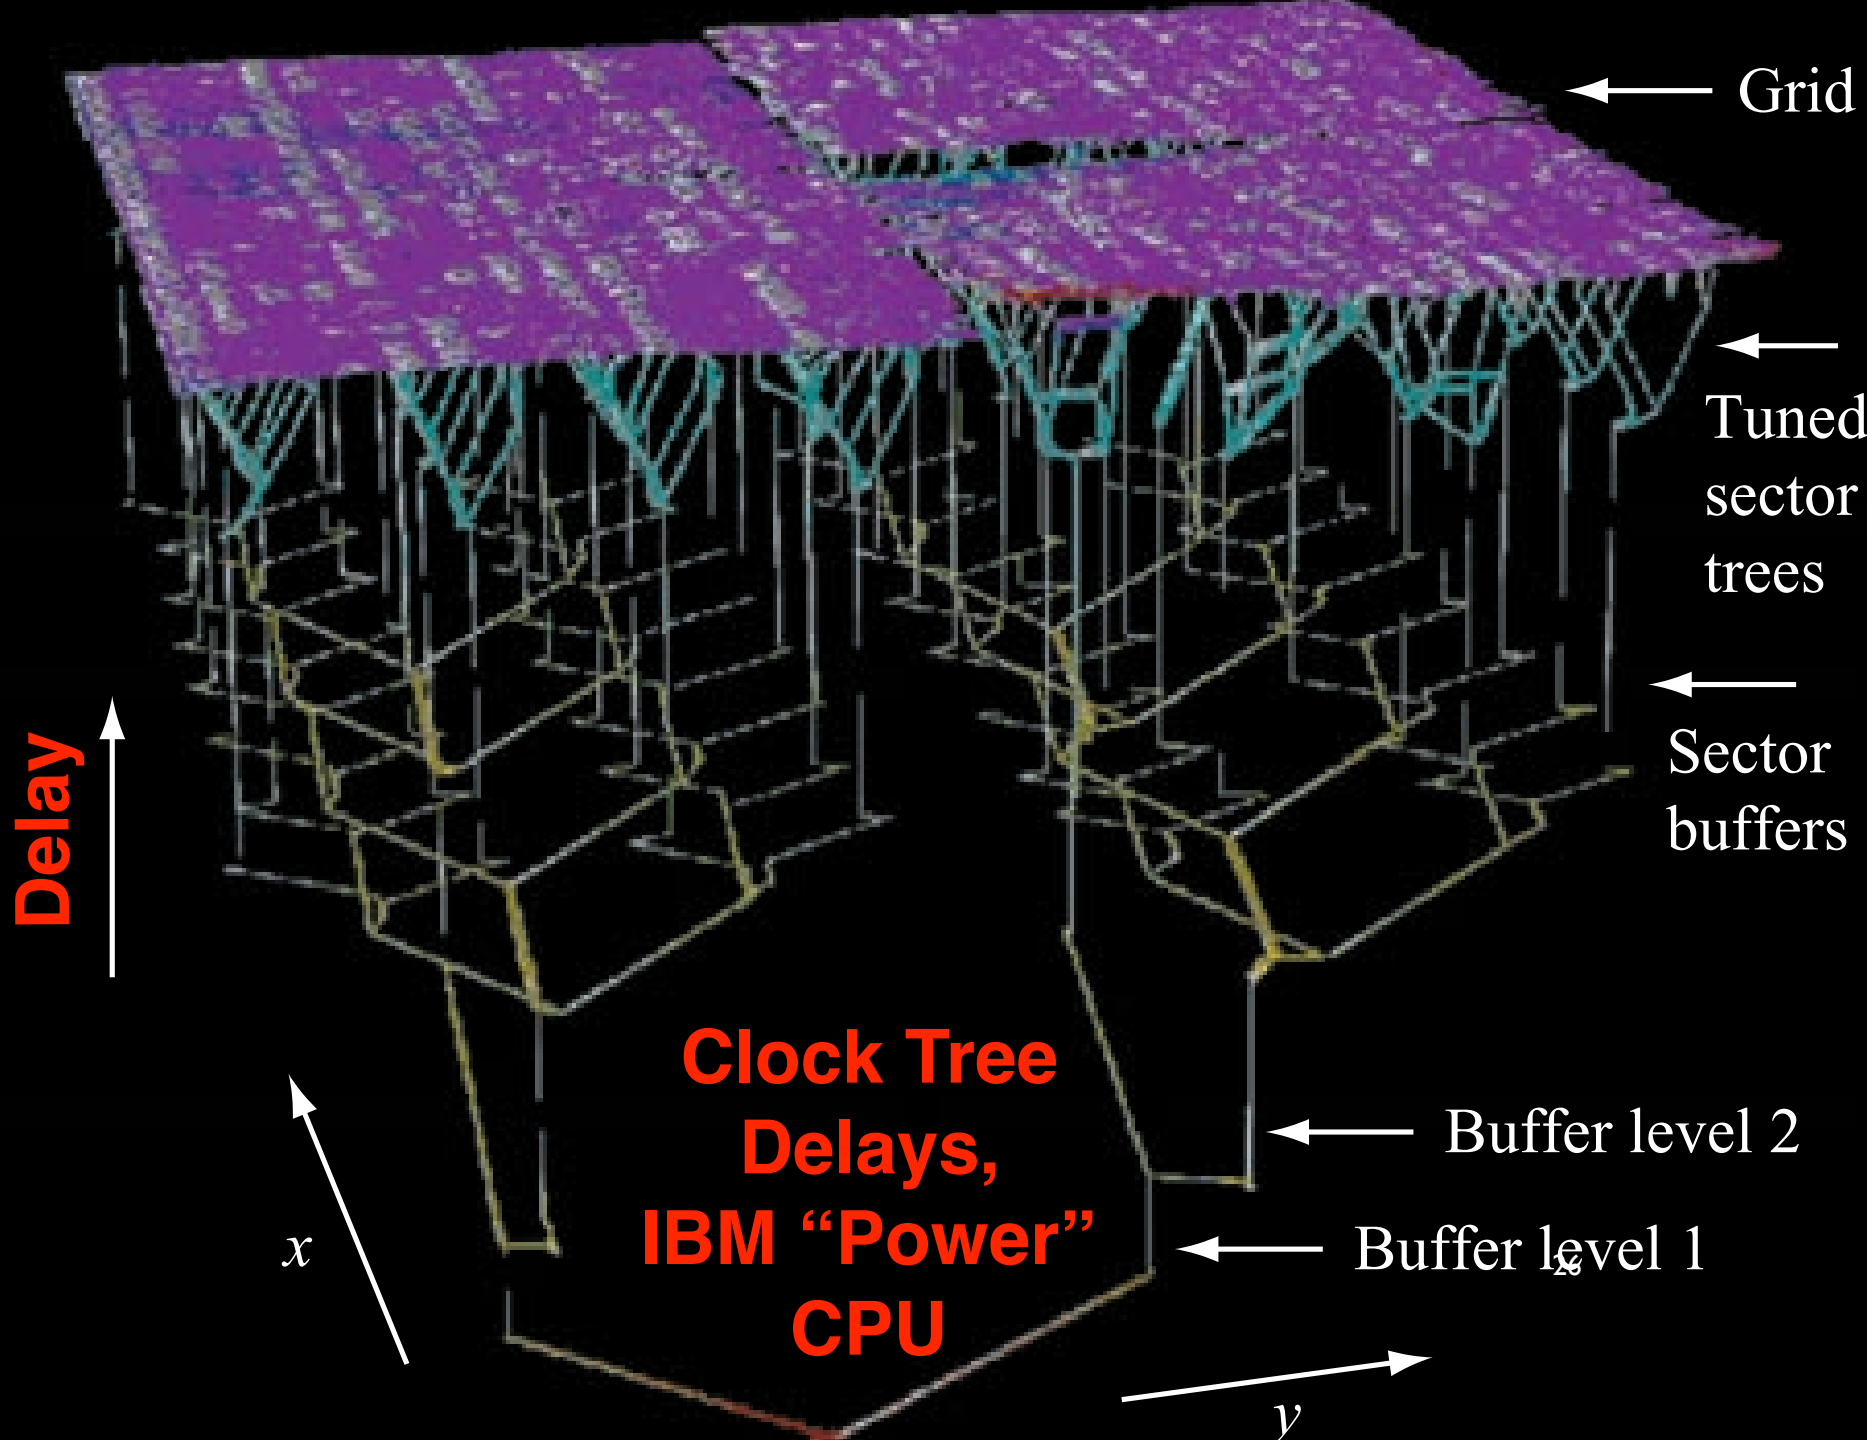



Equal wire length/number of buffers to get to every location

More realistic ASIC H-tree



[Restle98]







Clock Tree Delays, IBM Power



Chip Packaging


Chip Packaging




- ▶ Bond wires ($\sim 25\mu\text{m}$) are used to connect the package to the chip
- ▶ Pads are arranged in a frame around the chip
- ▶ Pads are relatively large
 - ▶ $\sim 100\mu\text{m}$ in $0.25\mu\text{m}$ technology, with $100\mu\text{m}$ pitch
 - ▶ $60\mu\text{m} \times 80\mu\text{m}$ at $80\mu\text{m}$ pitch in 45nm
- ▶ Many chips are 'pad limited'

Pad Frame



Layout



Die Photo



Chip Packaging





(b) Ball grid array packaging

- ▶ An alternative is 'flip-chip':
 - ▶ Pads are distributed around the chip
 - ▶ The solder balls are placed on pads
 - ▶ The chip is 'flipped' onto the package
 - ▶ Pads still large
 - ▶ But can have many more of them

Bonding Pad Design

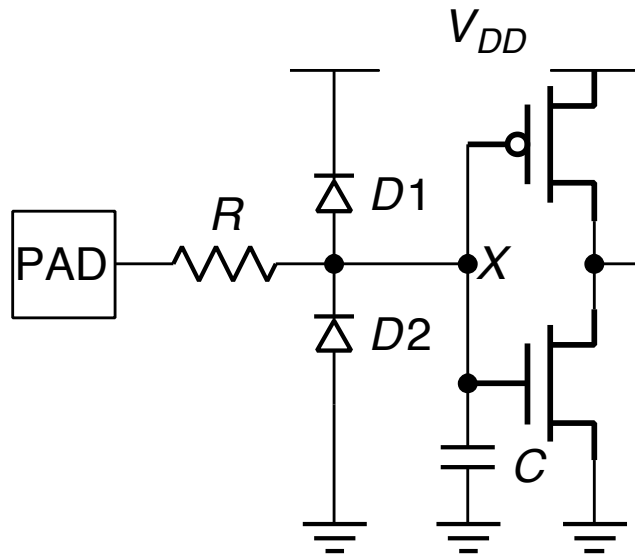
Bonding Pad





ESD Protection

- ▶ When a chip is connected to a board or otherwise handled, there is unknown (potentially large) static voltage difference (a few kV)
- ▶ Equalizing potentials requires (large) charge flow through the pads
- ▶ Diodes sink this charge into the substrate – need guard rings to pick it up.


Pads + ESD Protection



Diode



When Things Go Bad



[Maxim]



Power Distribution

Power Supply Distribution Issues


- ▶ **IR drops**
 - ▶ Voltage drops due to resistance in power wires
 - ▶ *Slower circuits, false switching*
- ▶ **Metal Migration (electromigration)**
 - ▶ *chip failures*
- ▶ **Inductive Effects**
 - ▶ bounce and oscillations on power nodes

All effects are helped by *shorter thicker wires*. Modern processes have special thick metal layers dedicated to power distribution. Area pads help keep connections to package short and distance from pad to circuit short.

Power Delivery

- ▶ Achieving good reliable power delivery requires a lot of resources:
 - ▶ ~70% of package pins just for power
 - ▶ Top 2-3 (thick) metal layers

All effects are helped by *shorter thicker wires*.
Modern processes have special thick metal layers dedicated to power distribution. Area pads help keep connections to package short and distance from pad to circuit short.




Chip Metal Layers



Electromigration


- ❑ If current density is too high - wires melts



- ❑ “On-chip wires: current limited to $\sim 1 \text{ mA}/\mu\text{m}$ for 5-7 year lifetime

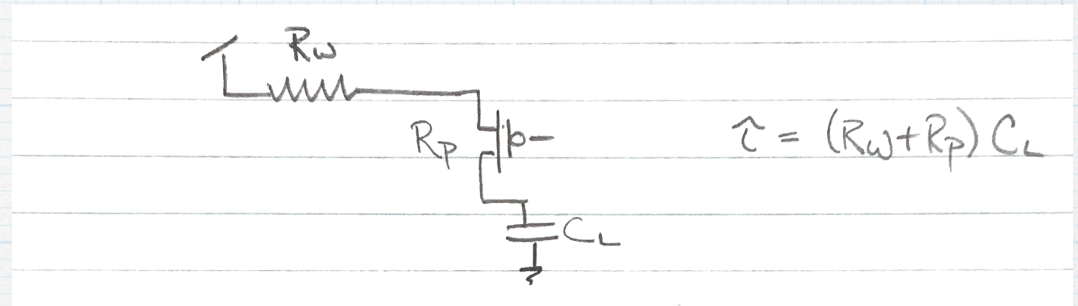
Power Supply Impedance (Z)

- ▶ Two principal elements increase Z:
 - ▶ Resistance of supply lines (IR drop)
 - ▶ Inductance of supply lines ($L \cdot di/dt$ drop)


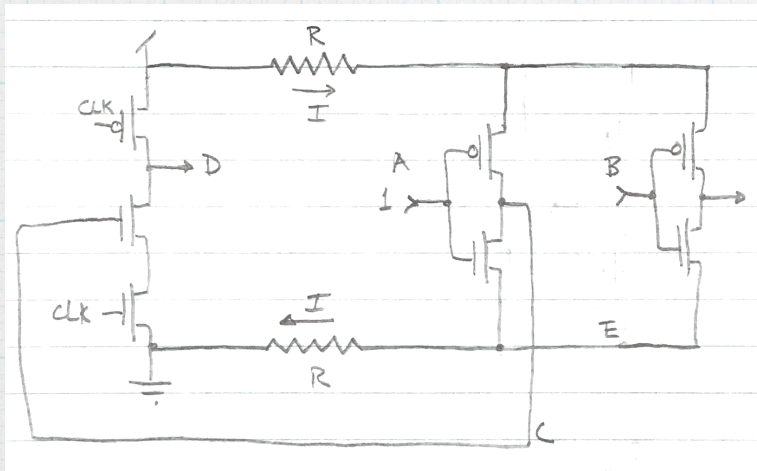


Power Supply Impedance

- ▶ IR voltage drop
- ▶ Slower circuit operation because of series resistance with transistors:



- ▶ IR drops generate “noise”:



Scaling and Supply Impedance

- ▶ Typical target for supply impedance is to get 5-10% voltage variation of nominal supply (e.g., 100mV for 1V supply)

- In traditional scaling V_{dd} drops while power stays constant.

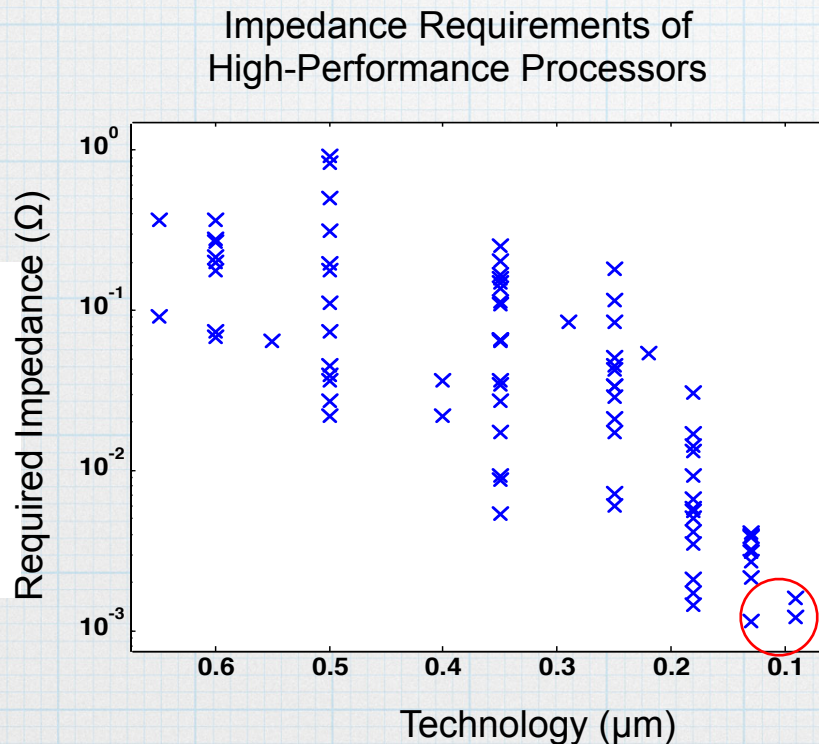
- This forced drastic drop in required supply impedance:

- $V_{dd} \downarrow, I_{dd} \uparrow \rightarrow |Z_{required}| \downarrow \downarrow$

- Extreme example:

- $V_{dd} = 1V, P=100W \Rightarrow I_{dd}=100A$

- For $\Delta V_{dd,max} = 100mV,$
 $Z_{dd,max} = 100mV/100A = 1m\Omega!$



IR Drop Example


- ▶ Intel Pentium 4: $\sim 103\text{W}$ at $\sim 1.275\text{V}$

$$I_{\text{dd}} = 81\text{Amps}$$

- ▶ For 10% IR drop, total distribution resistance must be less than $1.6\text{m}\Omega$
- ▶ On-chip wire $R \approx 20\text{m}\Omega/\text{sq.}$ (thick metal)
 - ▶ Can't meet R requirement even with multiple, complete layers dedicated to power
 - ▶ Main motivation for flip-chip packaging

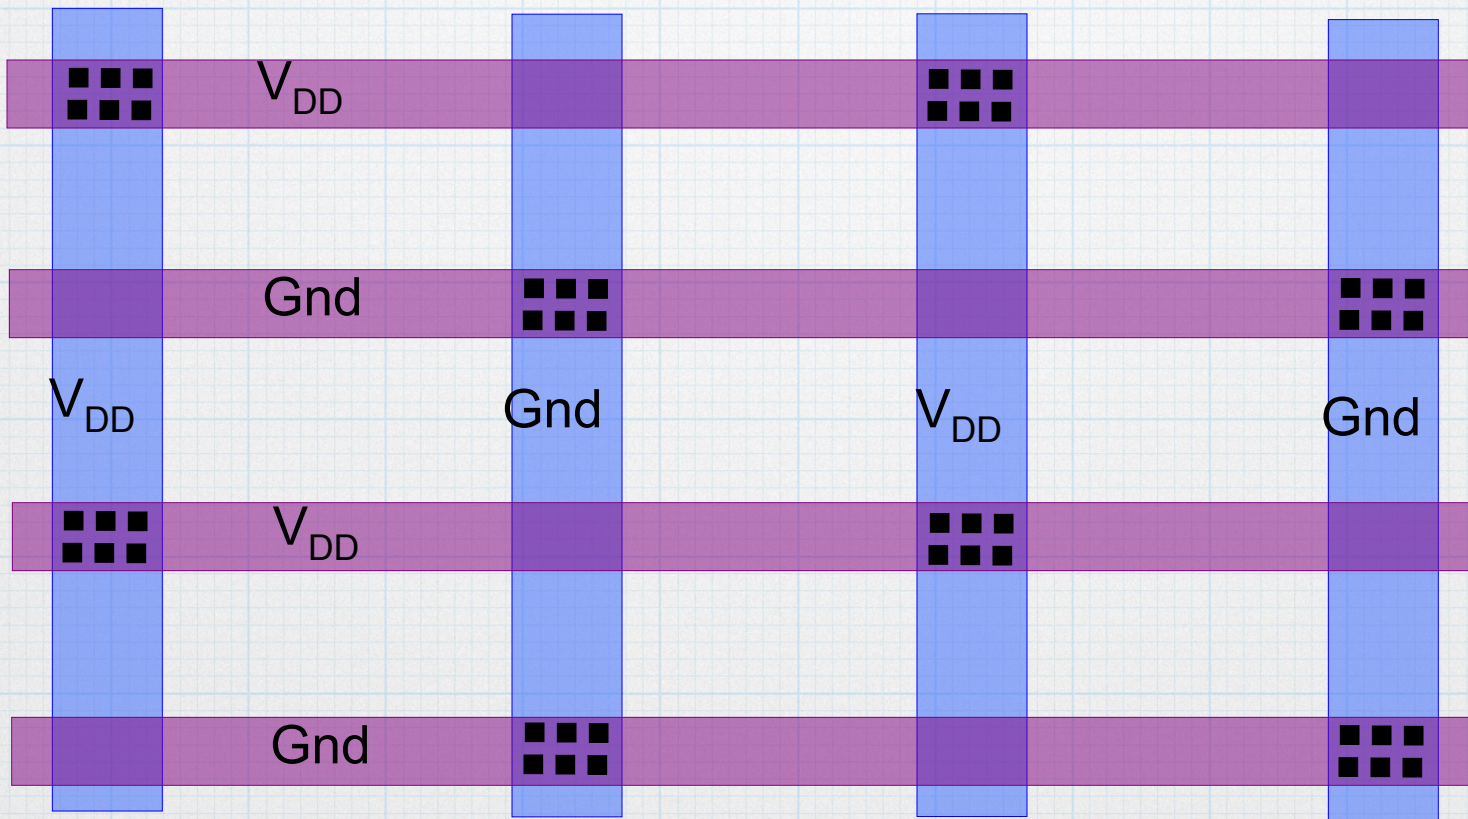
Layout Strategy

1. Keep distance from source of power/gnd as short as possible.
2. Use wide thick metal.
3. Isolate "noisy" sections.
4. Use multiple sources.



Popular On-Chip Power

- ▶ Power network usually follows pre-defined template (often referred to as “power grid”)

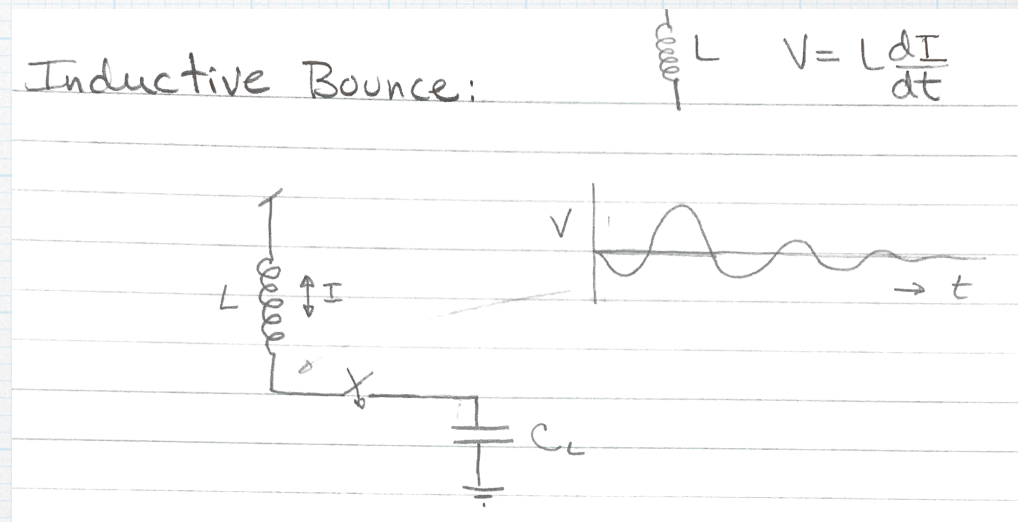


Inductive Bounce

Inductance in the power & ground paths results in voltage glitches (noise) on the Vdd & GND nodes.

On chip L value of wires is small => usually not significant except:

1. Very large currents: clock drivers, off-chip drivers
2. Package pins, bonding wires (1nH/mm)



Package pins can have from 2 - 40nH of inductance, depending on package type


Strategy:

1. Use multiple bonding pads (wiress) for Vdd and GND
2. Use on -chip by pass capacitors

Pin Inductance

- ▶ Major source of inductance is through the bonding pin connections to the chip package.
- ▶ C4 bump inductance is 25pH
- ▶ wire-bond inductance of 1nH/mm

$$V = L \cdot \frac{dI}{dt}$$



□ Example:


- Processor transient current is 50A in 20ps from 1V supply
- How many C4 bumps do we need to get supply noise spike of less than 10%?
- With wirebonds, how many wirebonds are needed?

Pin Inductance Example

- ▶ Processor transient current is 50A in 20ps from 1V supply:

$$V = L \cdot \frac{dI}{dt} \quad L = V \cdot \frac{dt}{dI} = 0.1V \cdot \frac{20ps}{50A} = 0.04pH$$

- ▶ C4 bump inductance is 25pH
- ▶ How many C4 bumps do we need to get supply noise spike of less than 10%?




Inductors in parallel 1/add

$$\frac{25pH}{0.04pH} = 625$$



$$\frac{1000pH}{0.04pH} > 25K!$$

On-chip Decoupling Capacitors Help with Inductive and Resistive Effects


- ▶ When transistors switch - current is drawn from C_D rather than through package pins and bonding wires - smooths our dI/dt .



- ▶ Distributed bypass capacitors also smooth out noise from IR drops.




MIM Cap : density (20-30 ff/um²)



Decap cell in stdcell library added by tools.

Decoupling Capacitors On-chip and on-board



Decoupling capacitors are added:

- ❑ On the board (right under the supply pins)
- ❑ On the chip (under the supply straps, near large buffers)
- ❑ C_d helps avoid current rushing through supply wires
 - ❑ local store of charge
 - ❑ “smoothing filter” on supply voltage

Decoupling Capacitors

- ▶ Under the die

