EECS151/251A
Spring 2018
Digital Design and Integrated Circuits

Instructors:
John Wawrzynek and Nick Weaver

Lecture 25:
Clock and Power Distribution

Outline

- Clock non-idealities
- Clock Distribution
- Chip Packaging
- Power Distribution

Administrative

- MT2 (finally) graded

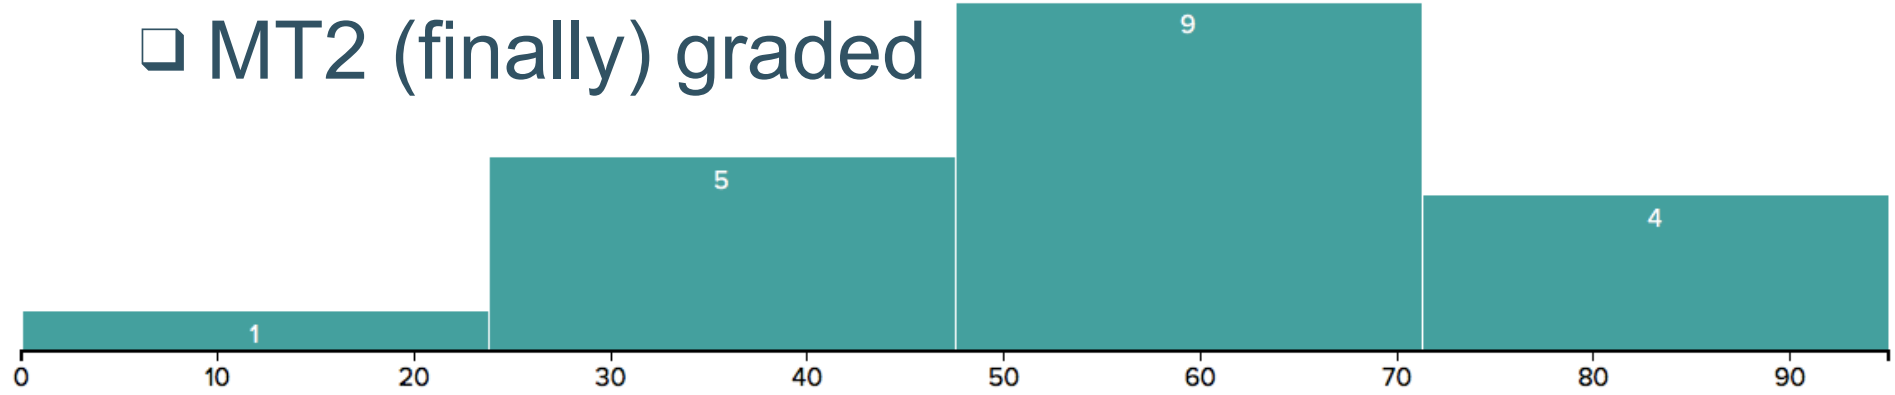

MINIMUM	MEDIAN	MAXIMUM	MEAN	STD DEV
17.0	59.0	86.0	57.45	16.92

- Final Project demo/interview
  - Friday 5/4 afternoon - signup for slot

Synchronous Timing - Review

Synchronous Timing

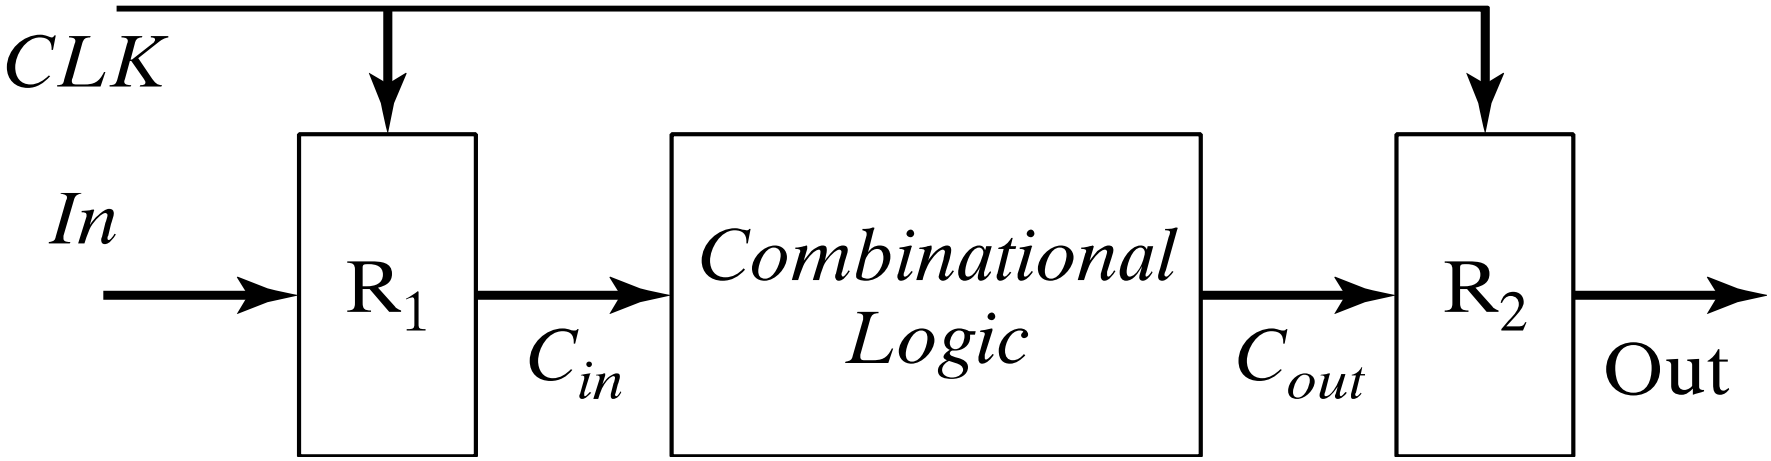

Latch Parameters

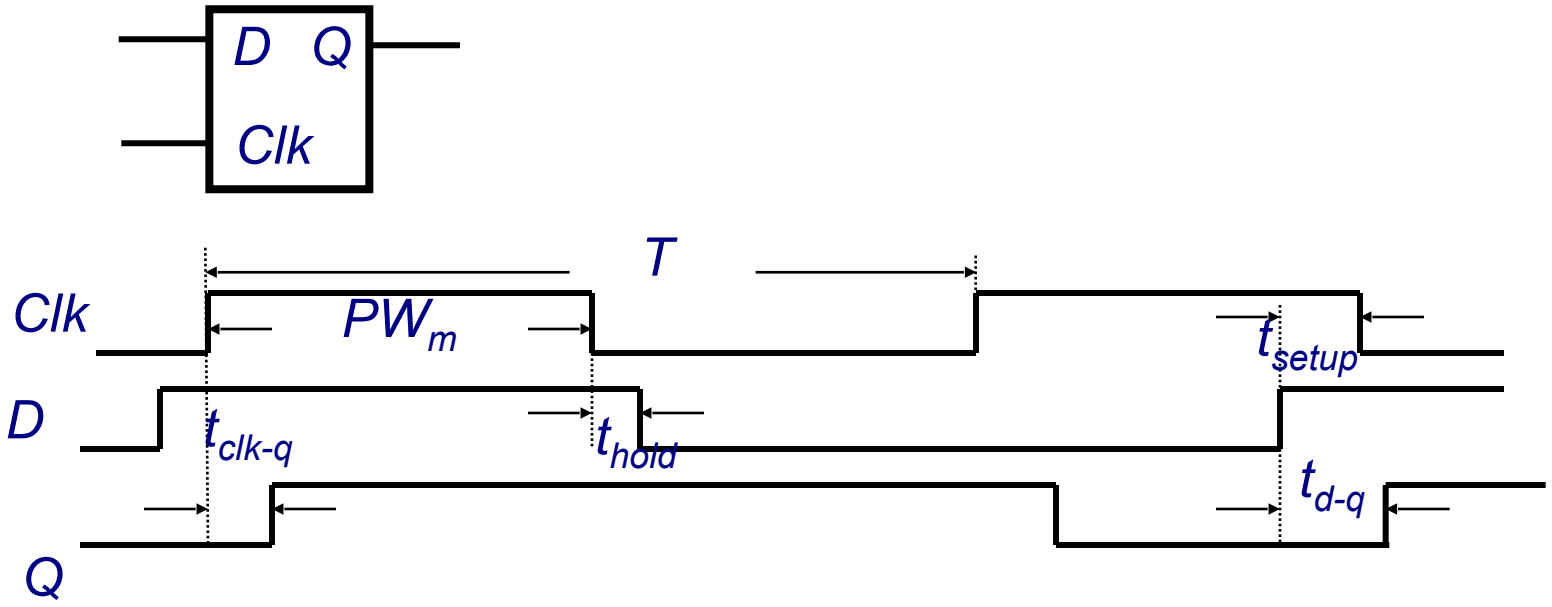

Delays can be different for rising and falling data transitions

Register Parameters

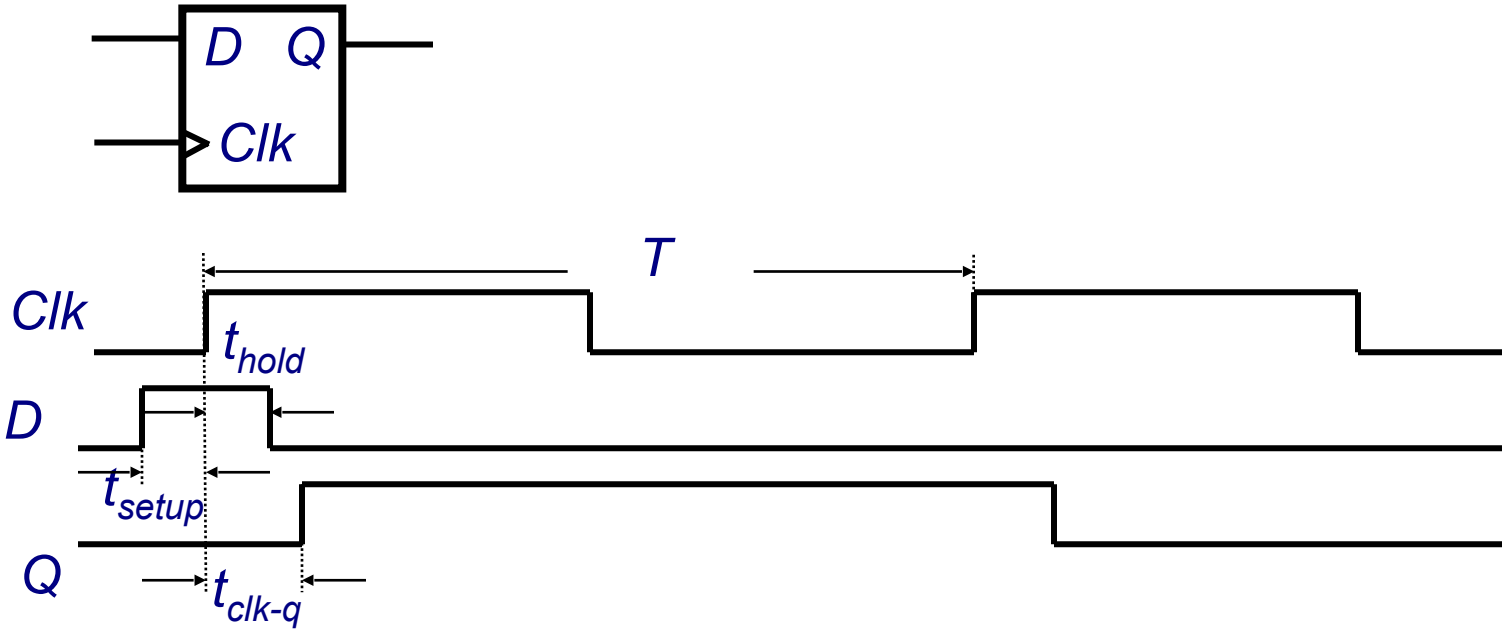

Delays can be different for rising and falling data transitions

Timing Constraints

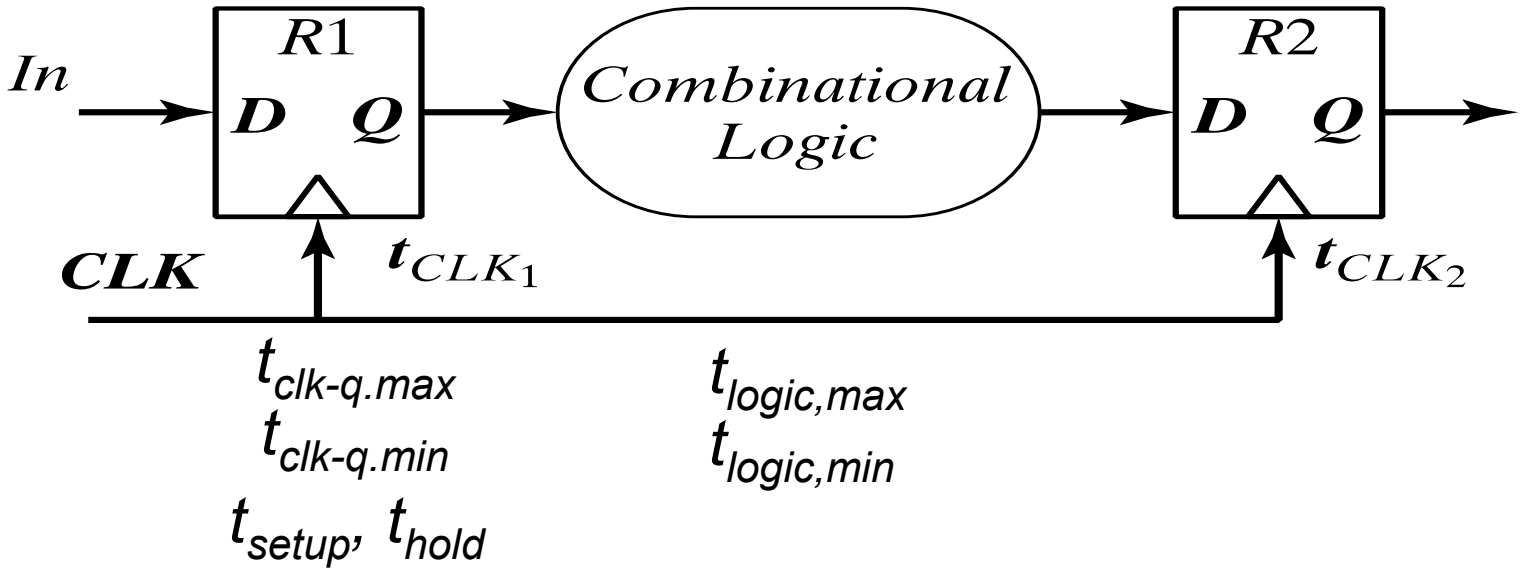

Timing Constraints

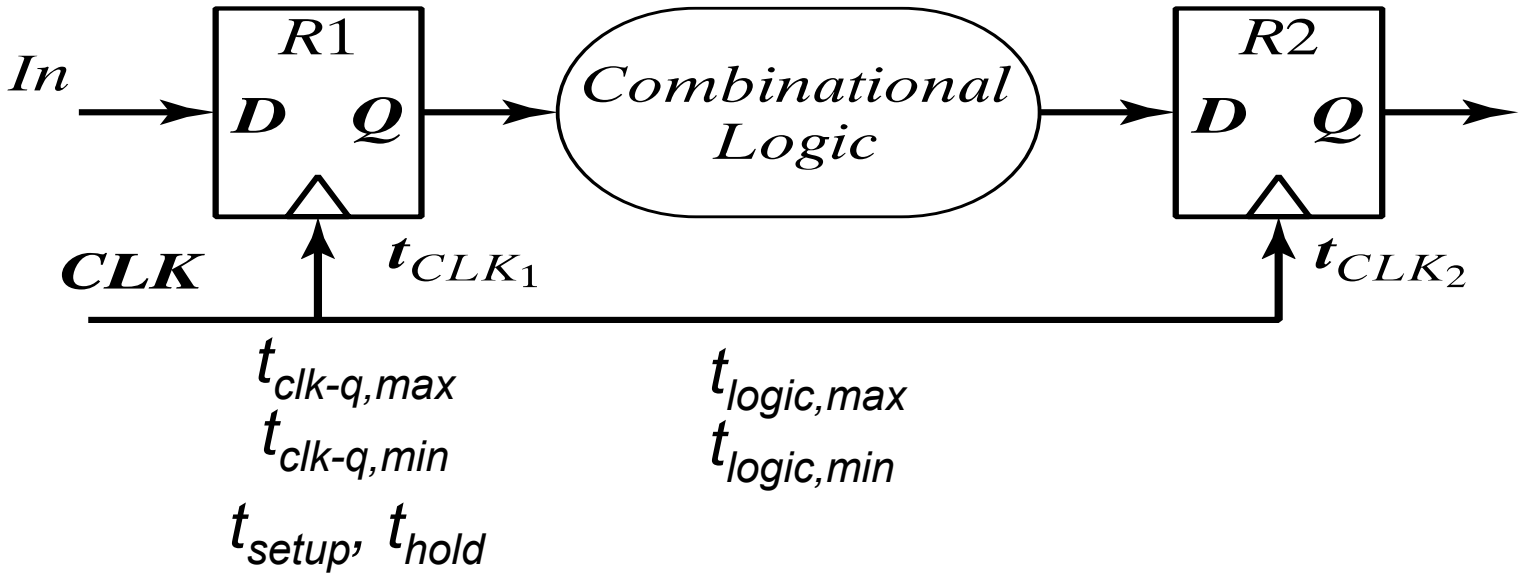

Cycle time (max): $T_{Clk} > t_{clk\text{-}q,max} + t_{logic,max} + t_{setup}$

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

Clock Nonidealities

Clock Nonidealities

- **Clock skew**
  - Time difference between the sink (receiving) and source (launching) edge
  - Spatial variation in temporally equivalent clock edges; deterministic + random, t_{SK}
- **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

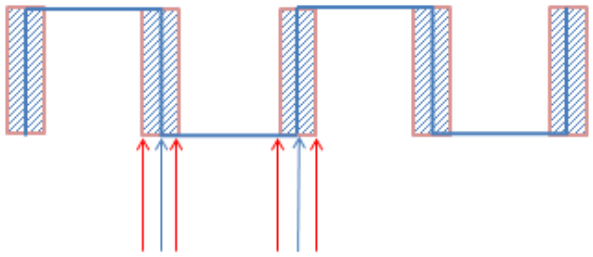

Clock Uncertainties

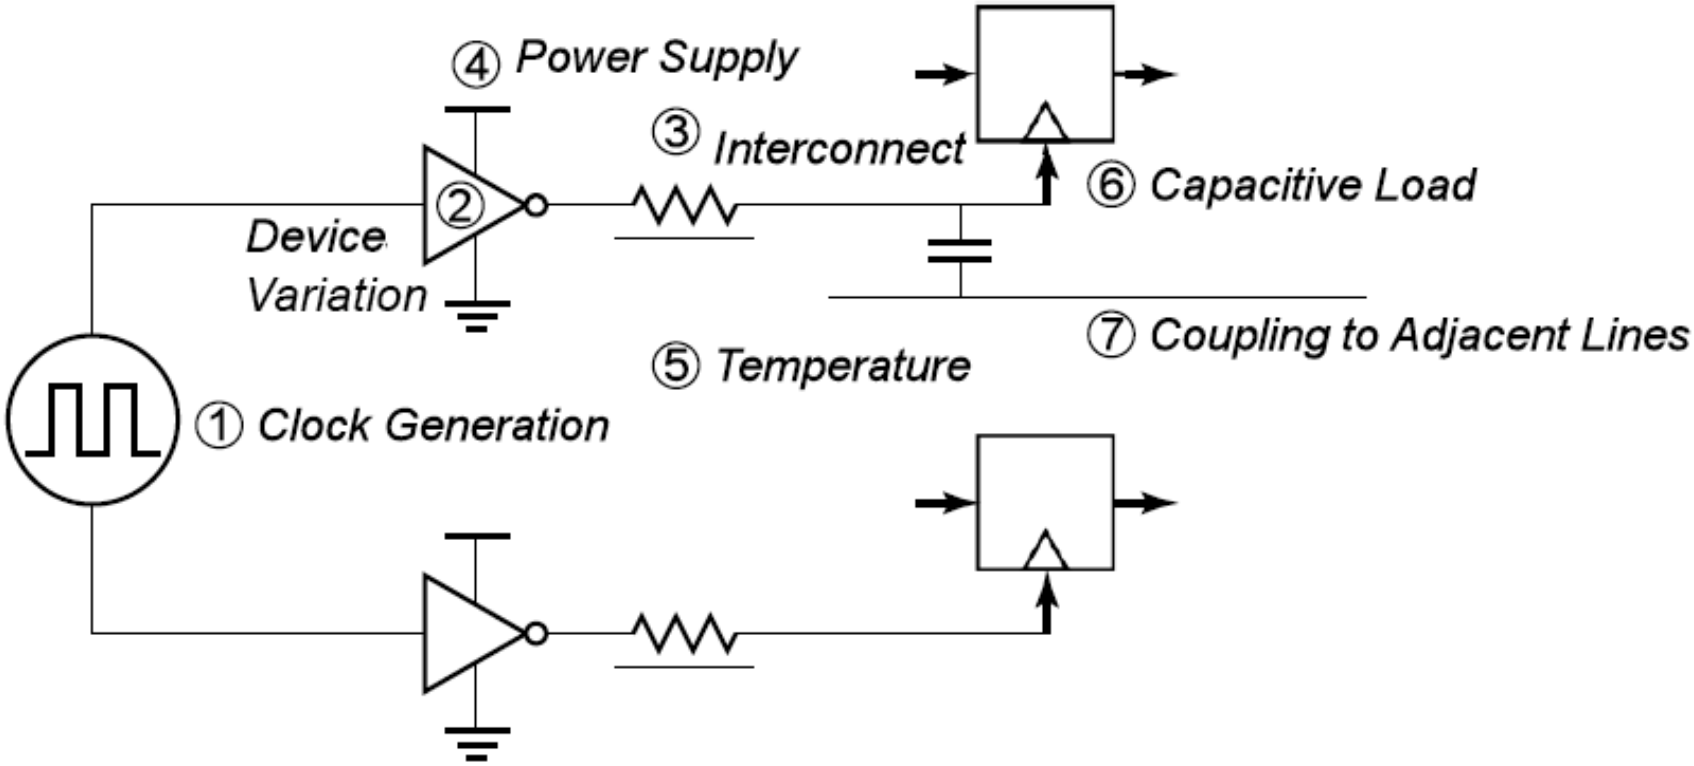

Sources of clock uncertainty

Clock Skew and Jitter

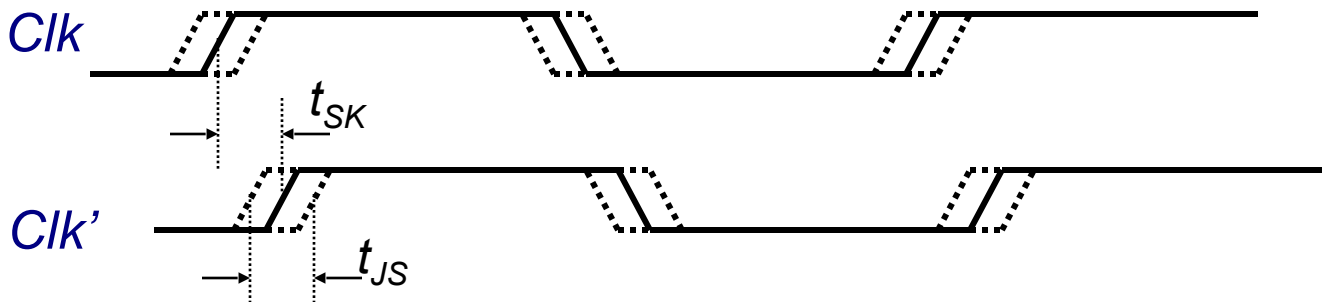

- Both skew and jitter affect the effective cycle time and the race margin

Positive Skew

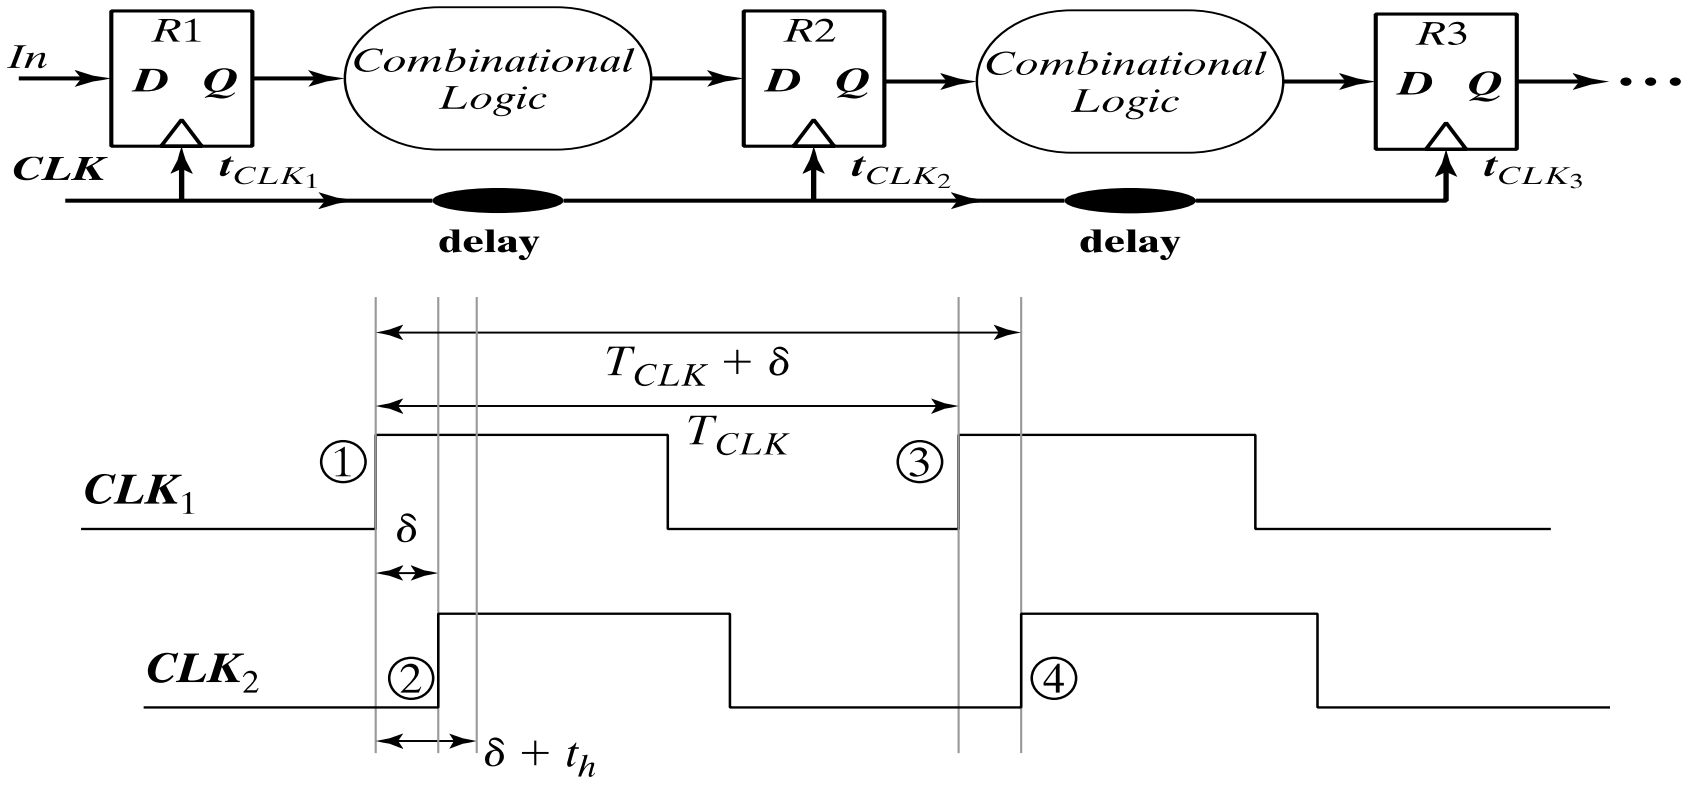

Launching edge arrives before the receiving edge

Negative Skew

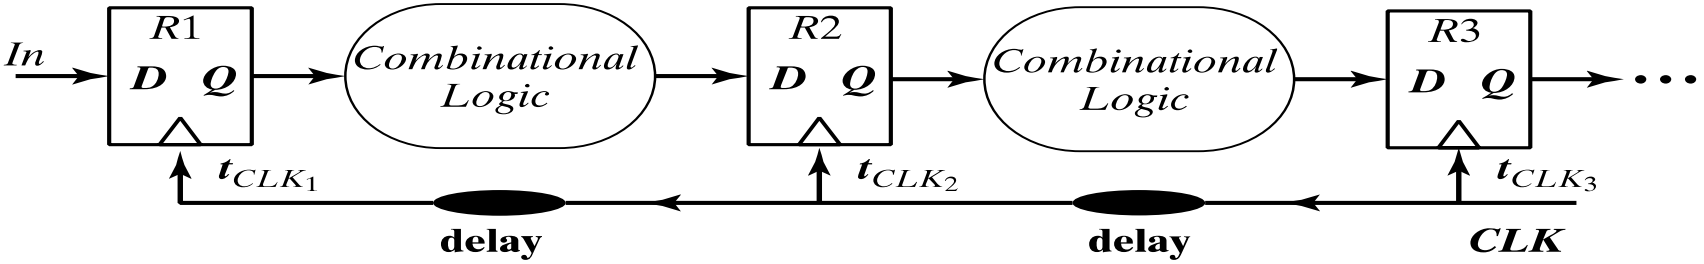

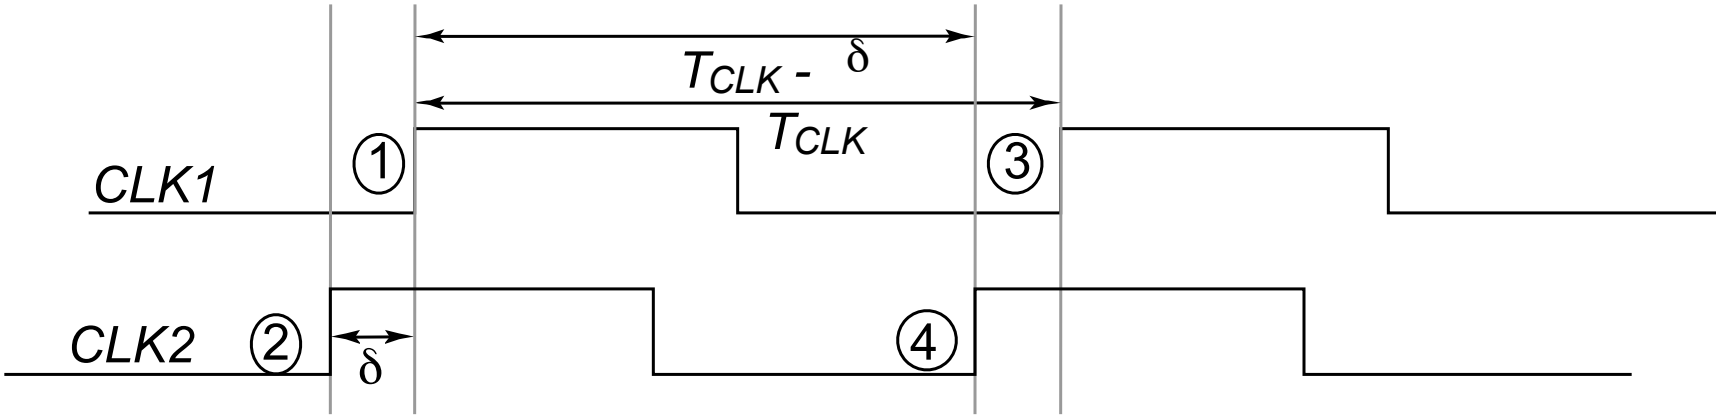

Receiving edge arrives before the launching edge

Timing Constraints

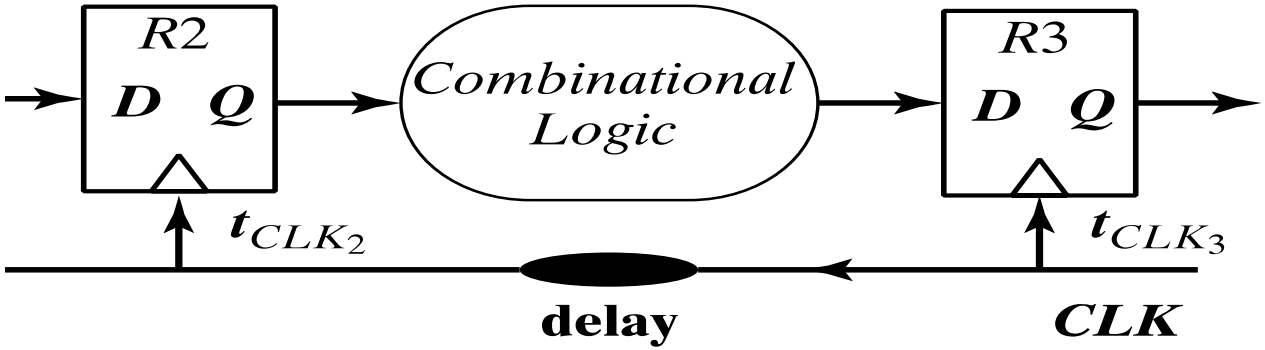

$t_{clk\text{-}q,max}$
$t_{clk\text{-}q,min}$
t_{setup}, t_{hold}

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

Minimum cycle time:

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

Timing Constraints

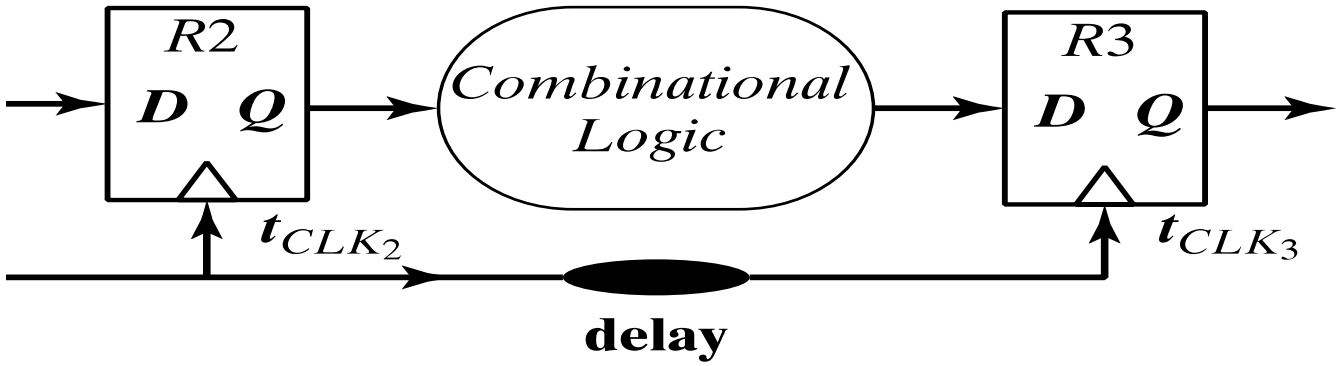

$t_{clk\text{-}q,max}$ $t_{logic,max}$

$t_{clk\text{-}q,min}$ $t_{logic,min}$

t_{setup}, t_{hold}

Hold time constraint:

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

Longest Logic Path in Edge-Triggered Systems

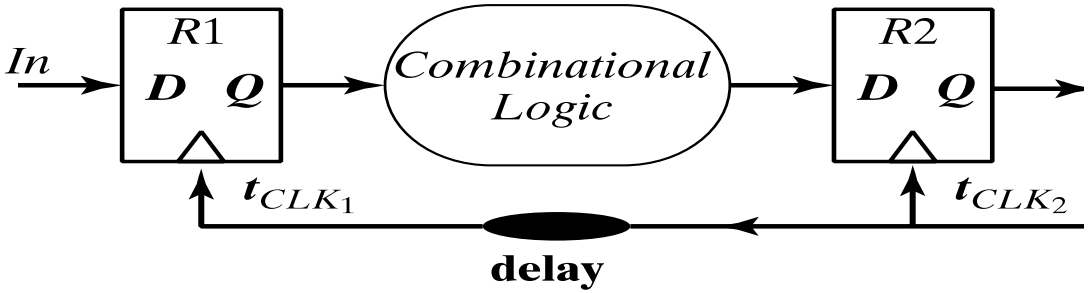

Latest point of launching

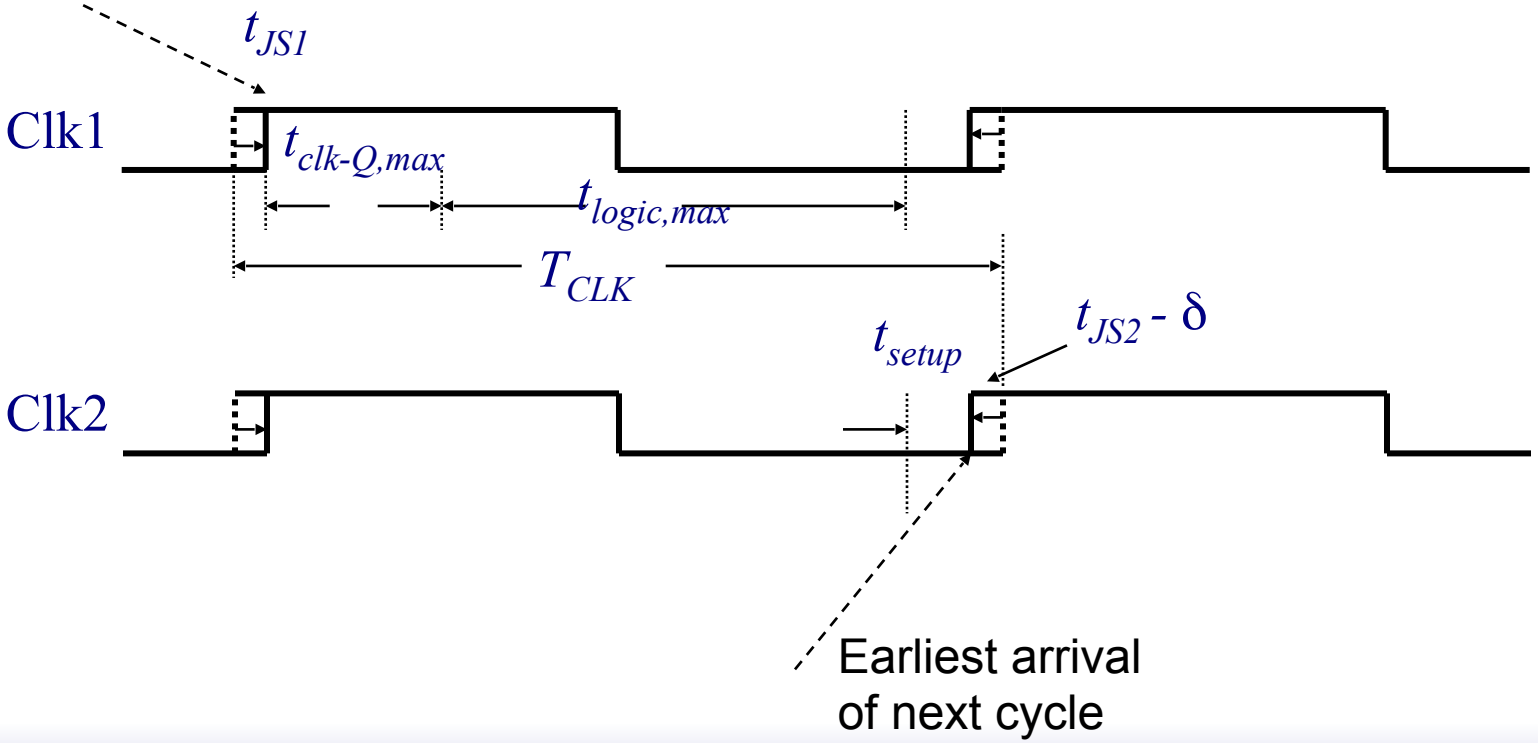

Earliest arrival of next cycle

Clock Constraints in Edge-Triggered Systems

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

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

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

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

Skew can be either positive or negative

Jitter t_{JS} usually expressed as peak-to-peak or N x rms value

Shortest Path

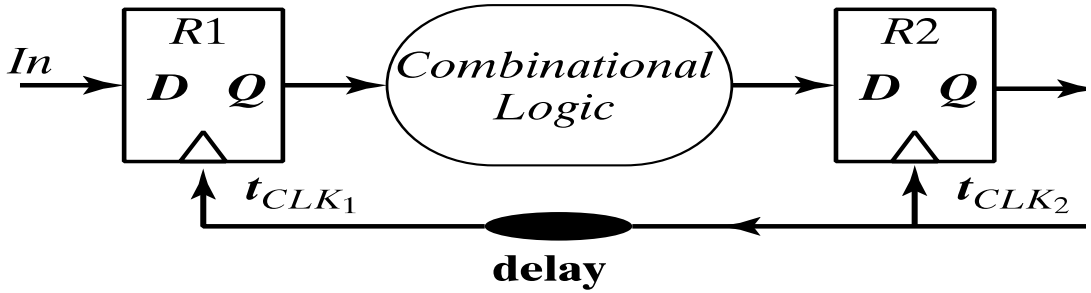

Earliest point of launching

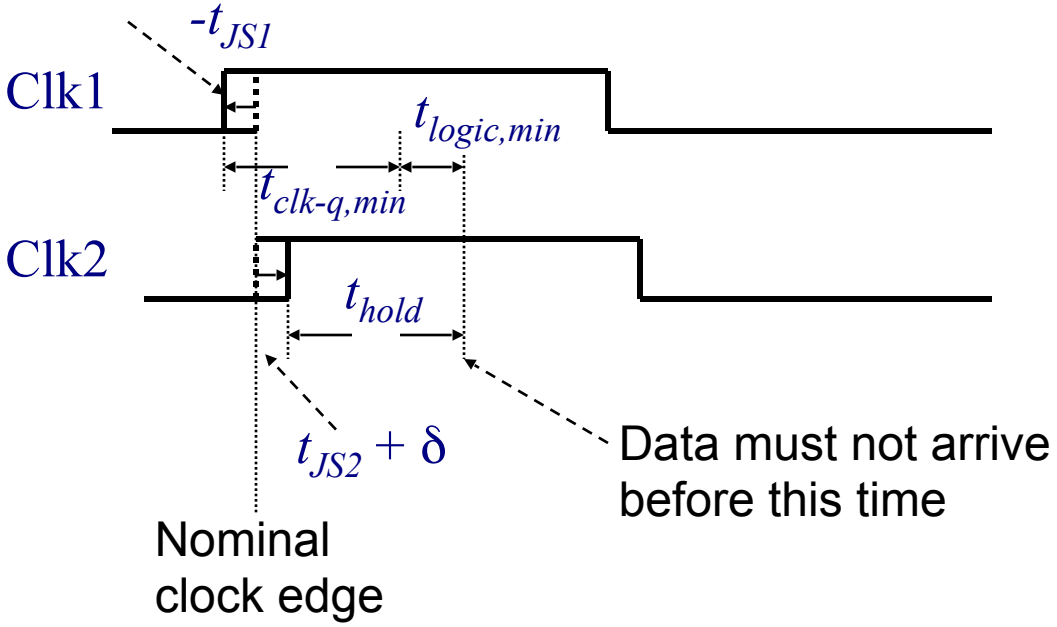

Clock Constraints in Edge-Triggered Systems

If launching edge is early and receiving edge is late:

$$t_{clk\text{-}q,min} + t_{logic,min} - t_{JS,1} > t_{hold} + t_{JS,2} + \delta$$

Minimum logic delay

$$t_{clk\text{-}q,min} + t_{logic,min} > t_{hold} + 2t_{JS} + \delta$$

(This assumes jitter at launching and receiving clocks are independent – which usually is not true)

Datapath with Feedback

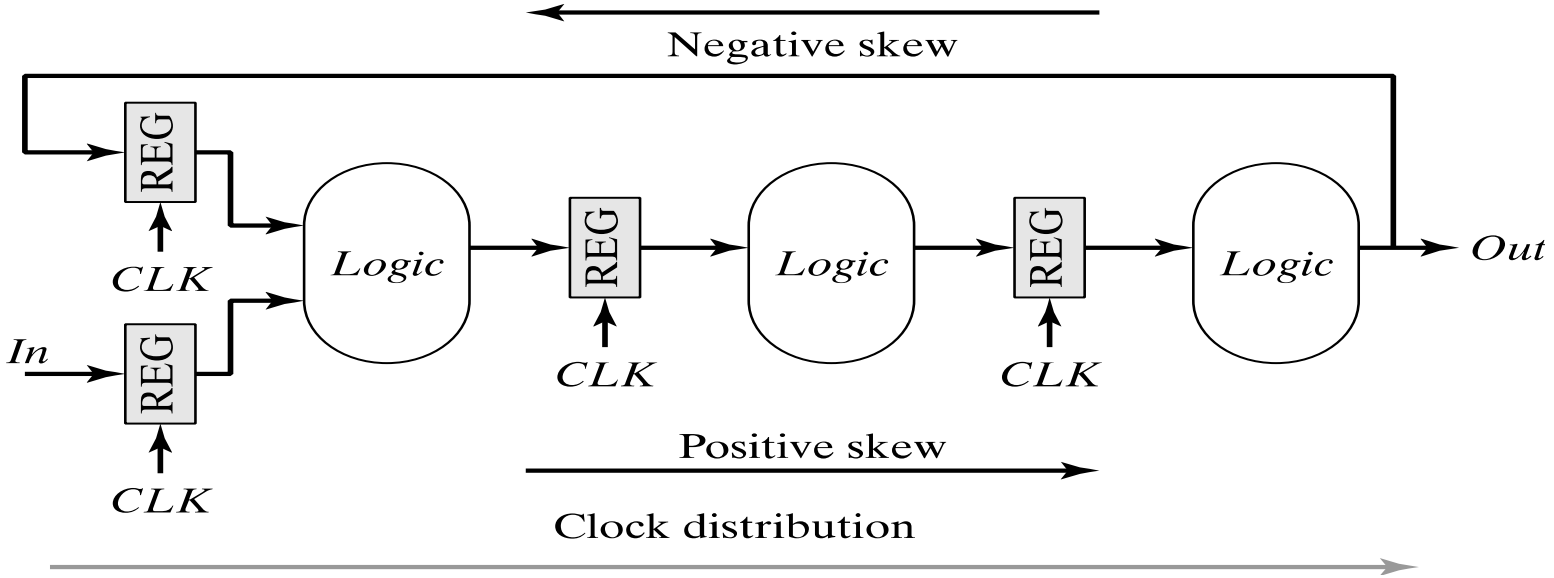

Clock Distribution

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?

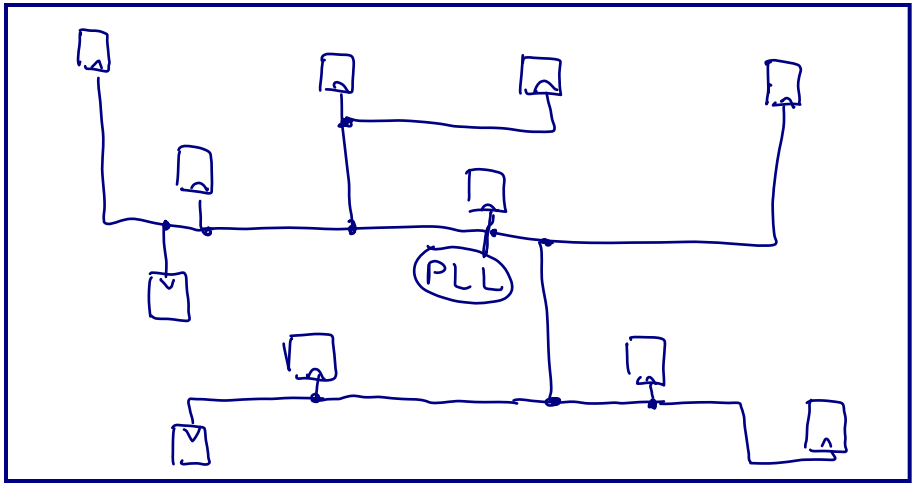

H-Tree

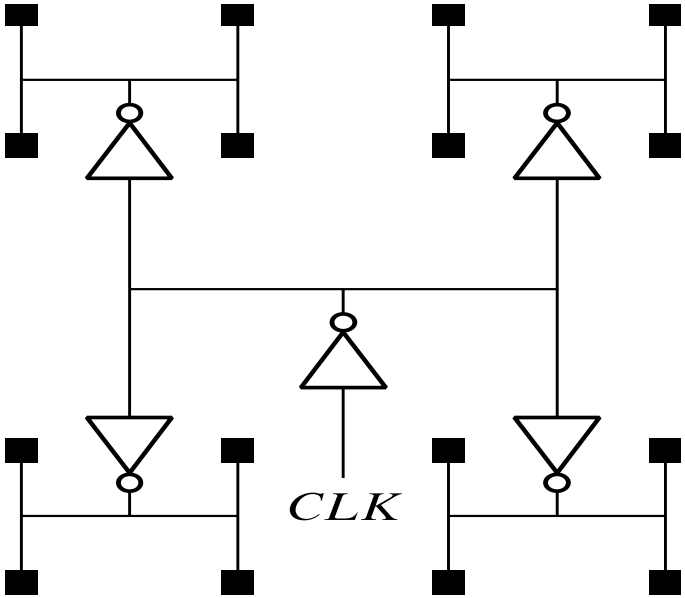

Equal wire length/number of buffers to get to every location

More realistic H-tree

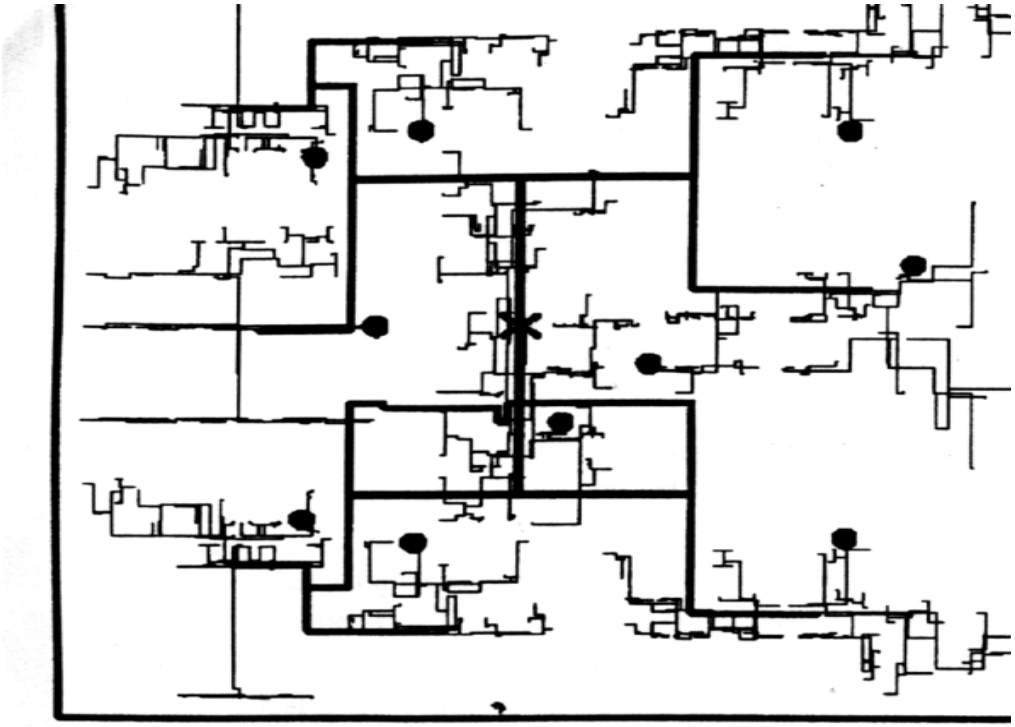

[Restle98]

Clock Grid

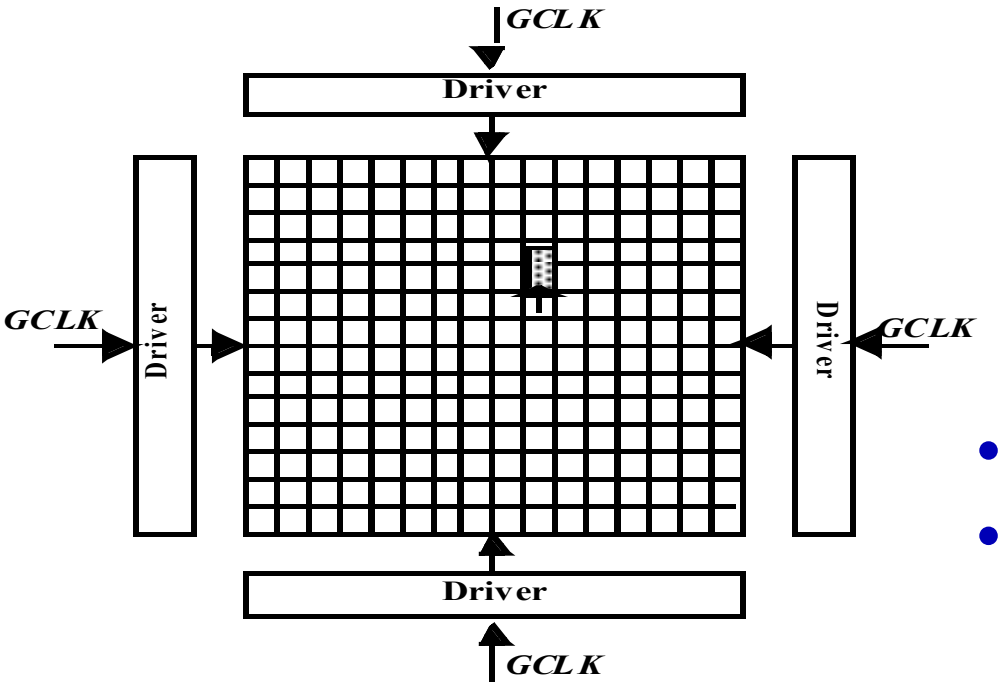

- No RC matching
- But huge power

Example: DEC Alpha 21164 (1995)

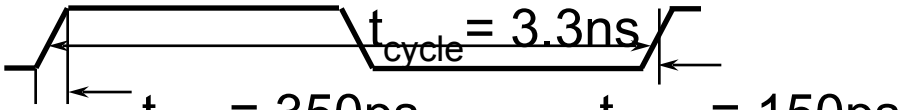

Clock waveform

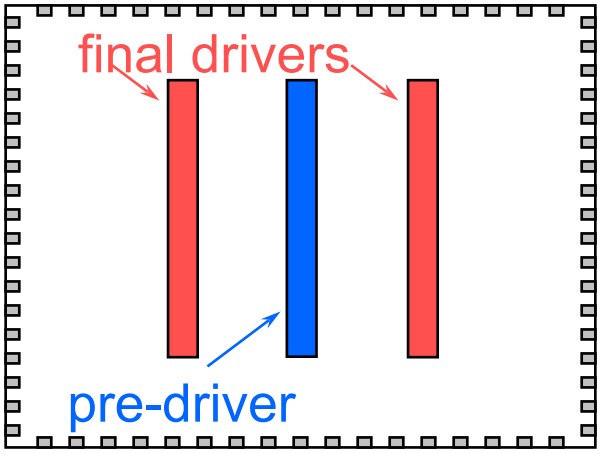

Location of clock driver on die

- 2 phase single wire clock, distributed globally
- 2 distributed driver channels
  - Reduced RC delay/skew
  - Improved thermal distribution
  - 3.75nF clock load, 20W power
  - 58 cm final driver width
- Local inverters for latching
- Conditional clocks in caches to reduce power

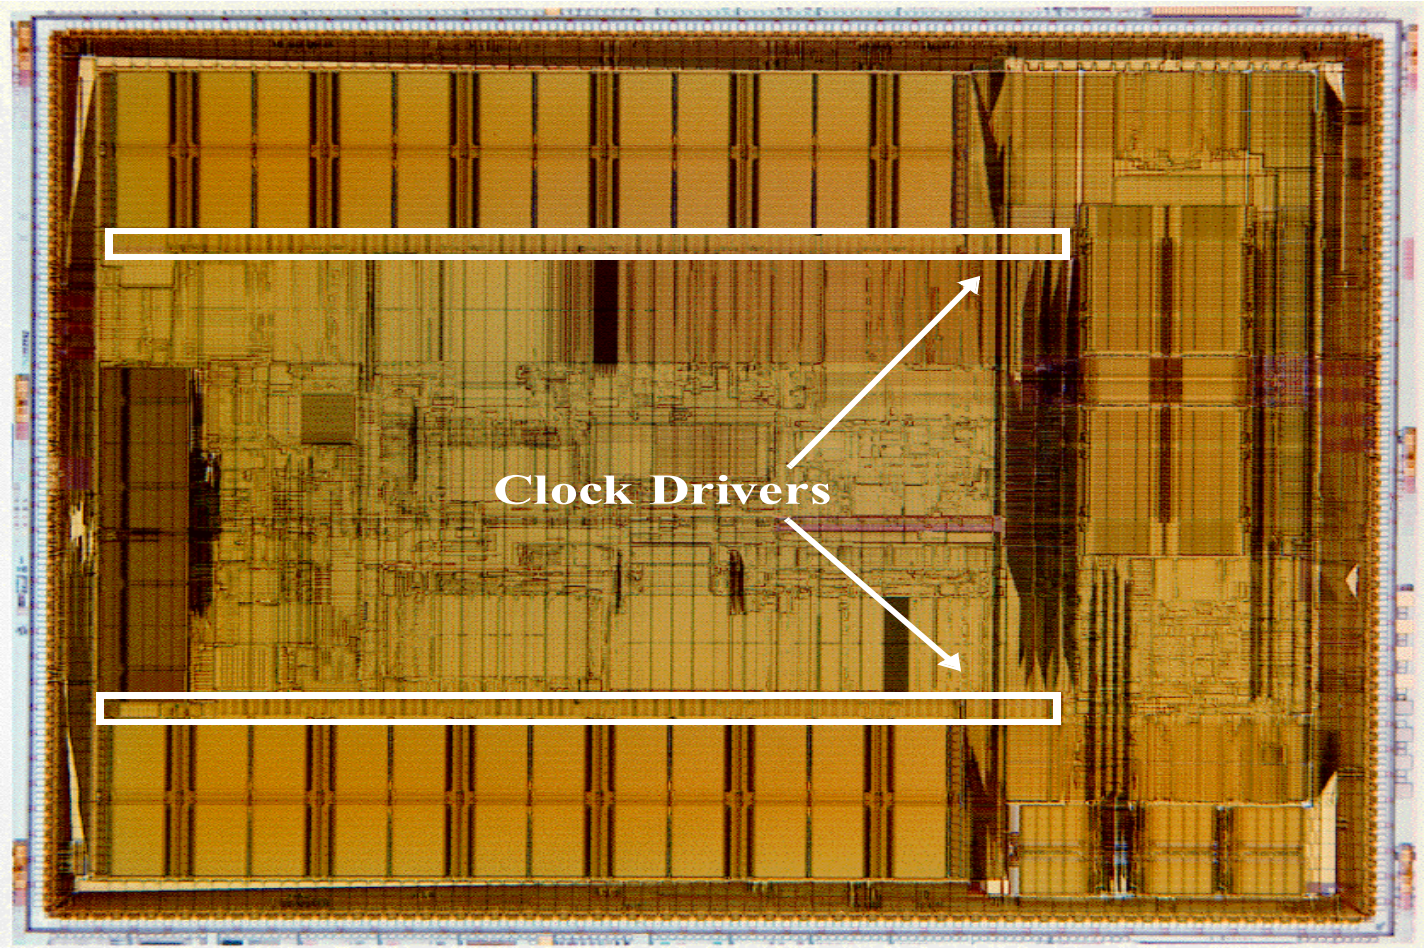
Clock Drivers

Clock Skew in Alpha Processor

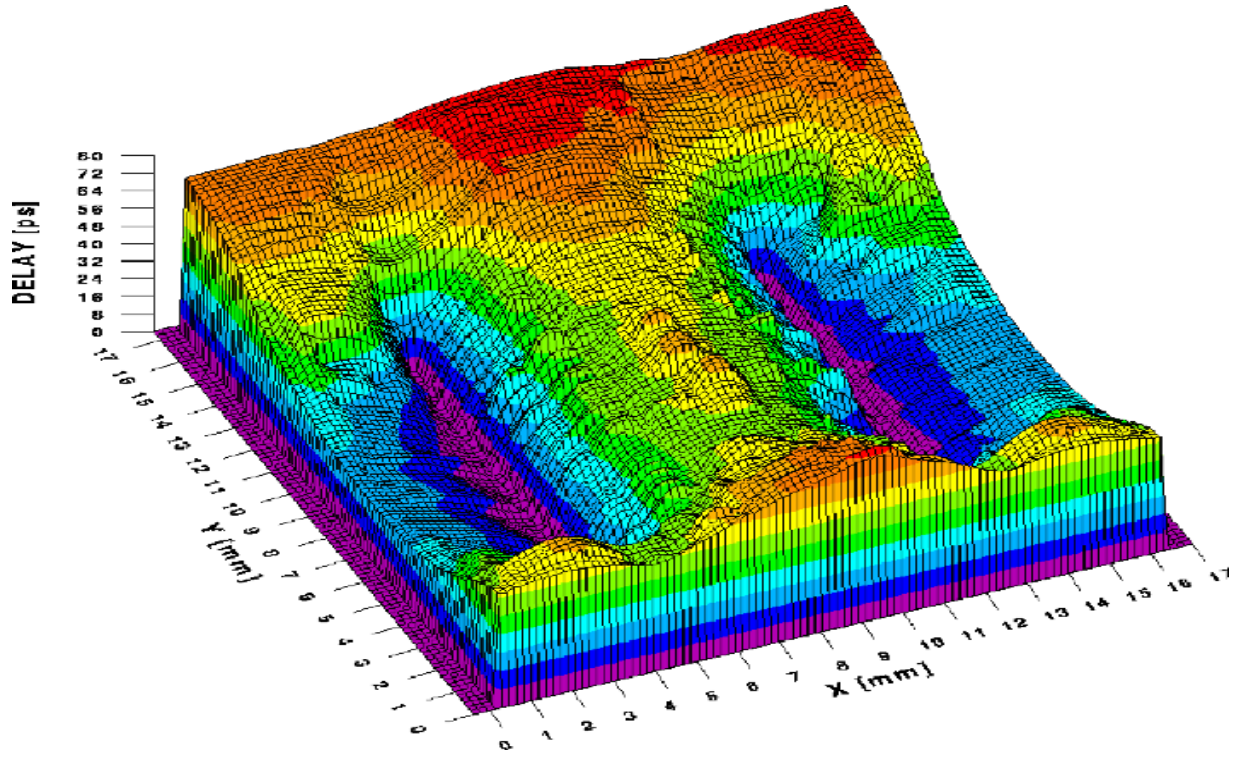

EV6 (Alpha 21264) Clocking
600 MHz – 0.35 micron CMOS

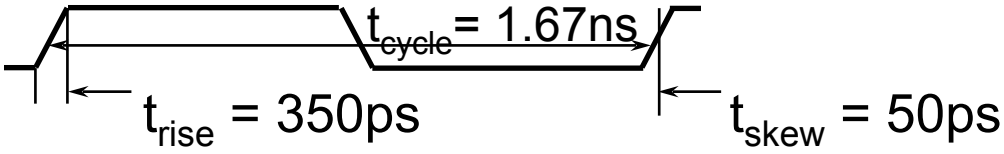

Global clock waveform

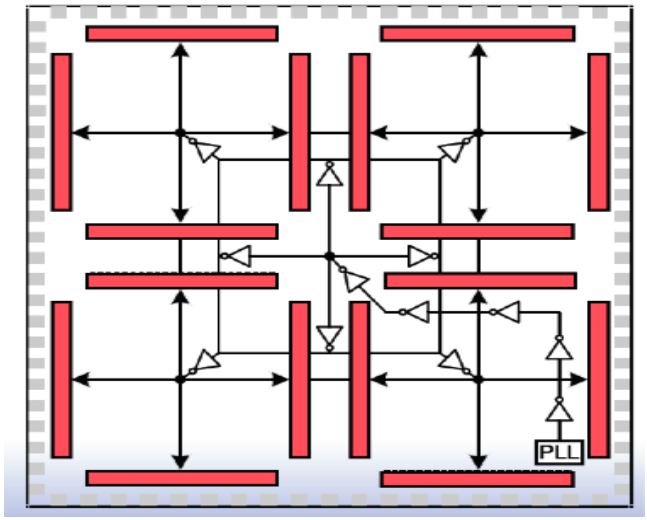

- 2 Phase, with multiple conditional buffered clocks
  - 2.8 nF clock load
  - 40 cm final driver width
- Local clocks can be gated “off” to save power
- Reduced load/skew
- Reduced thermal issues
- Multiple clocks complicate race checking

21264 Clocking

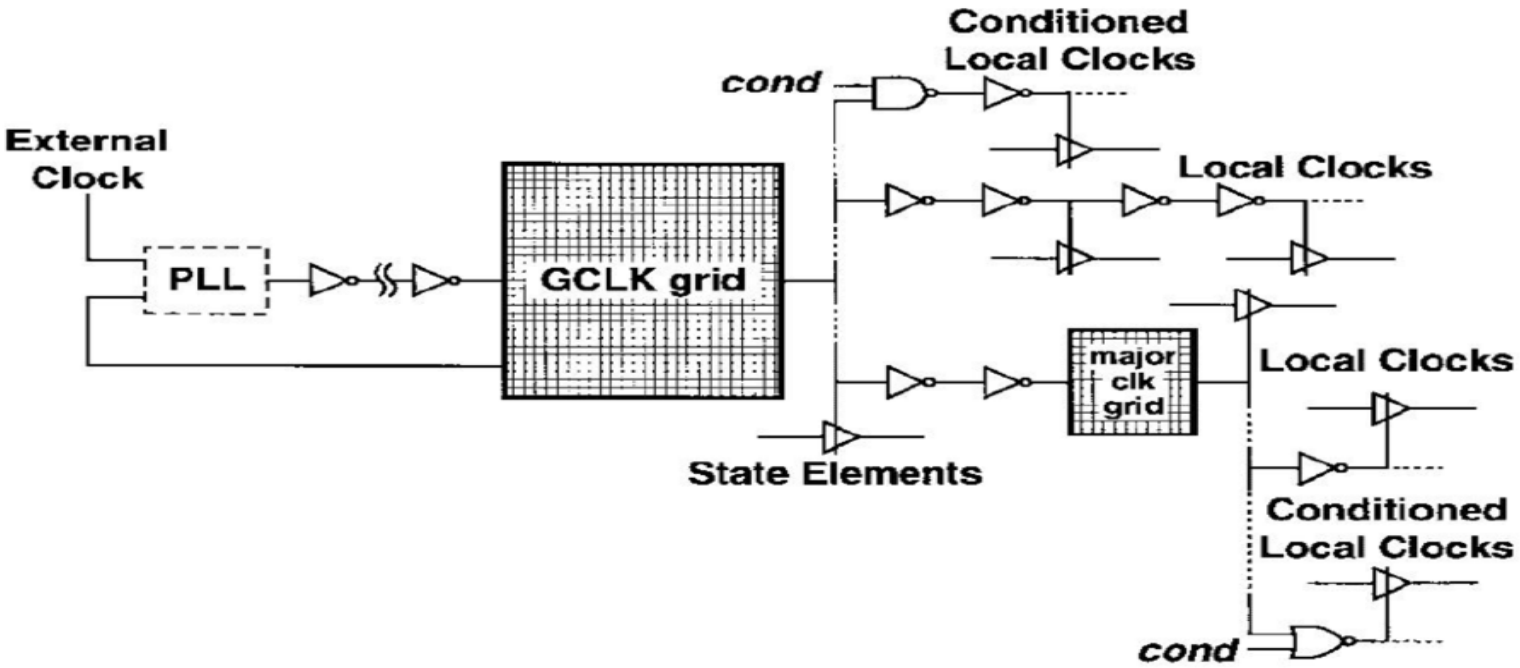

EV6 Clock Results

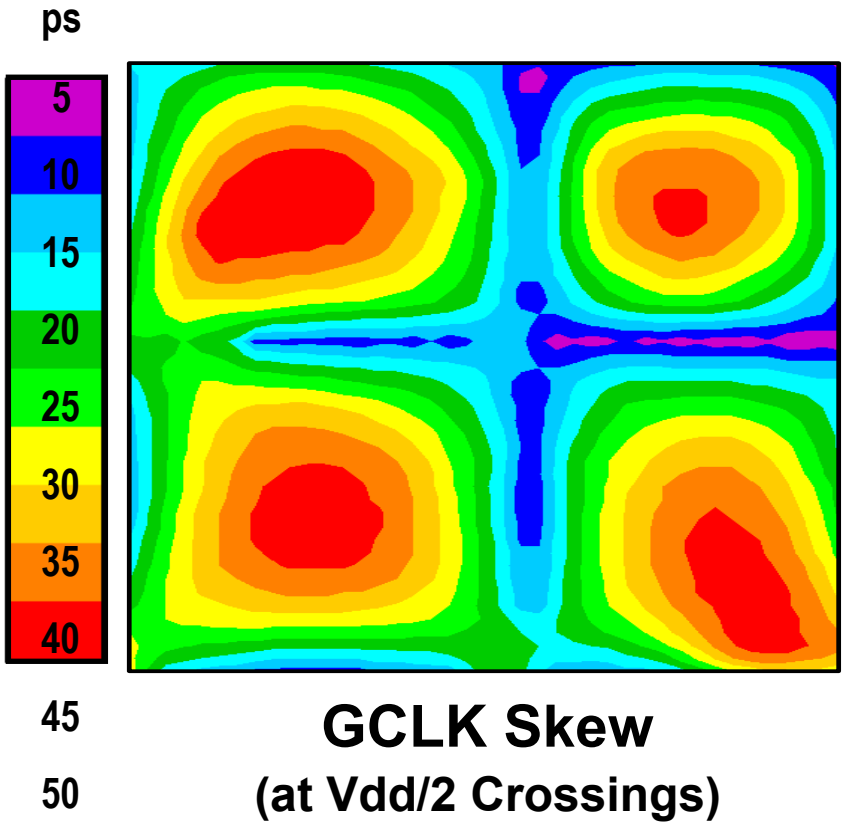

GCLK Skew
(at Vdd/2 Crossings)

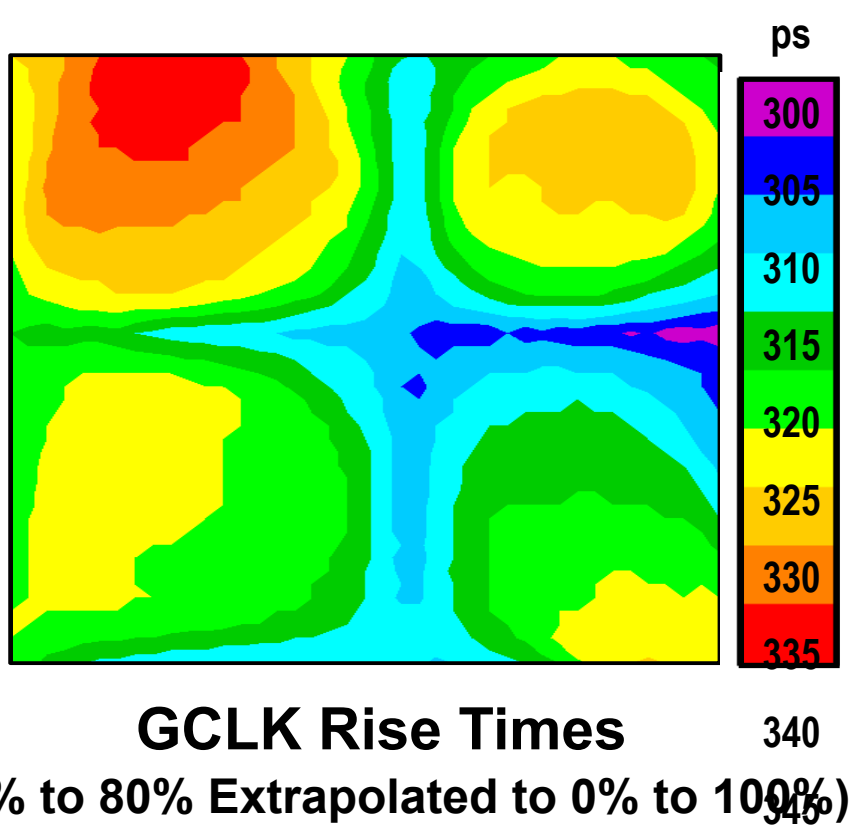

GCLK Rise Times
(20% to 80% Extrapolated to 0% to 100%)

EV7 Clock Hierarchy (2002)

Active Skew Management and Multiple Clock Domains

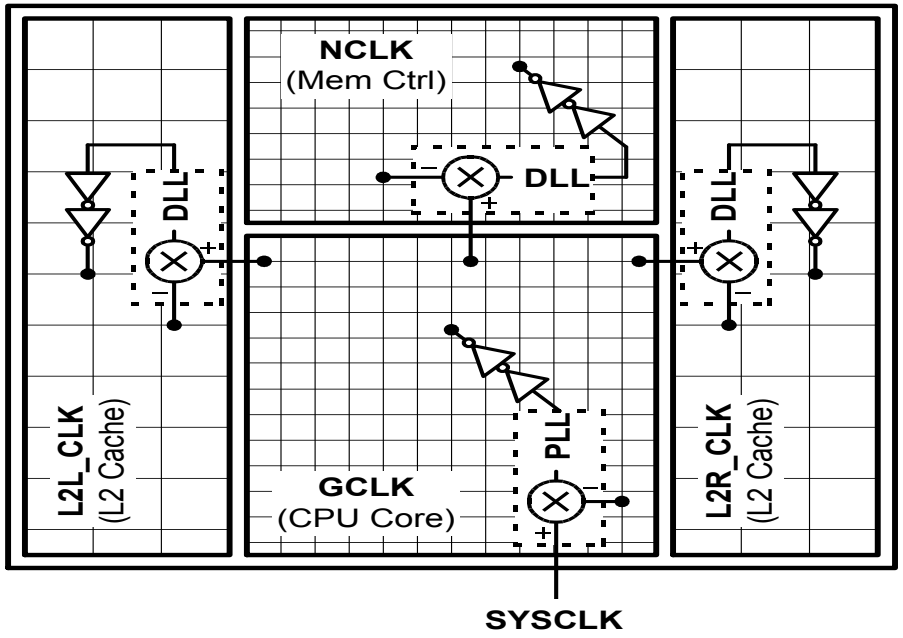

+ widely dispersed drivers

+ DLLs compensate static and low-frequency variation

+ divides design and verification effort

– DLL design and verification is added work

+ tailored clocks

Modern Processors

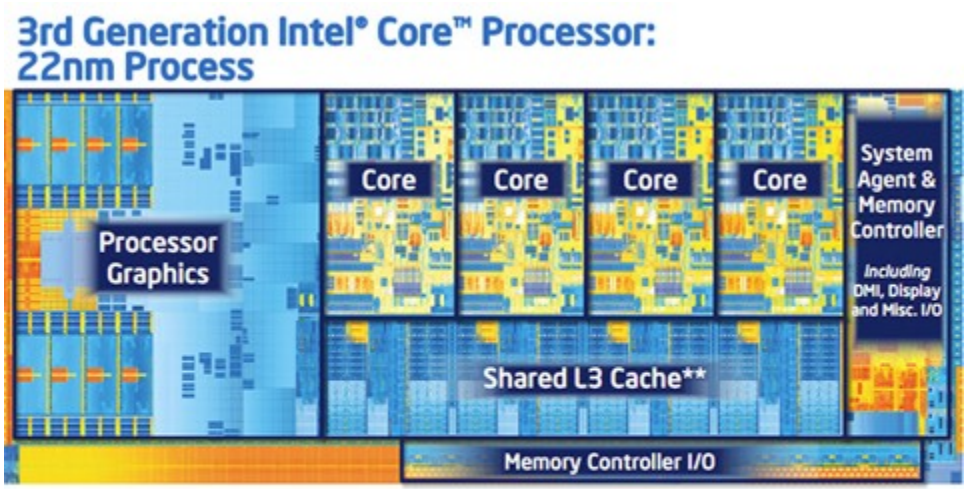

Apple A7

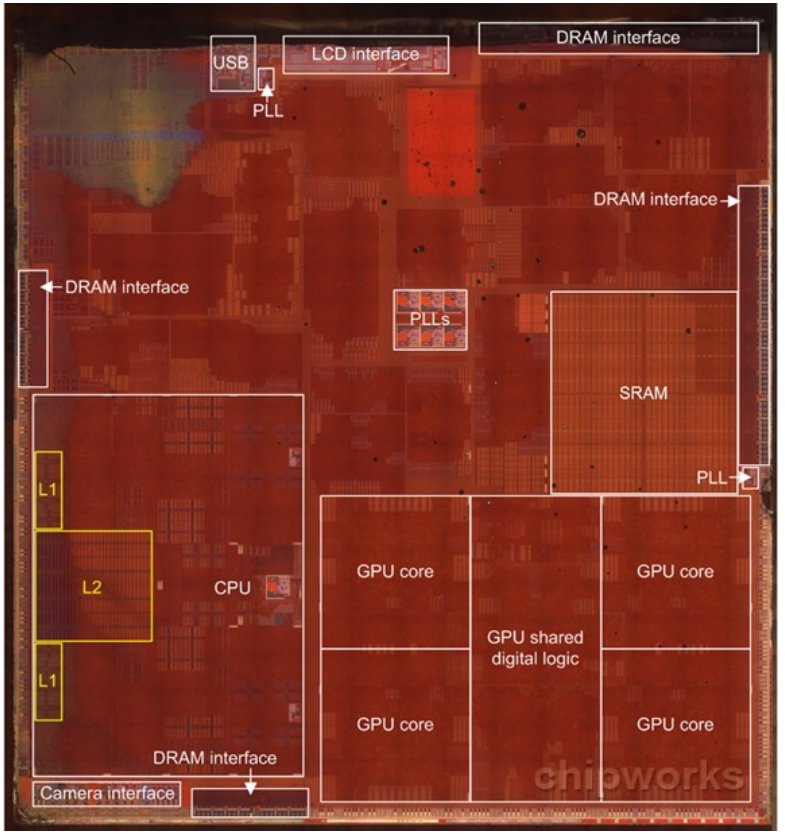

Clock Animations

- By Phillip Restle (IBM)

http://www.research.ibm.com/people/r/restle/resonate.html

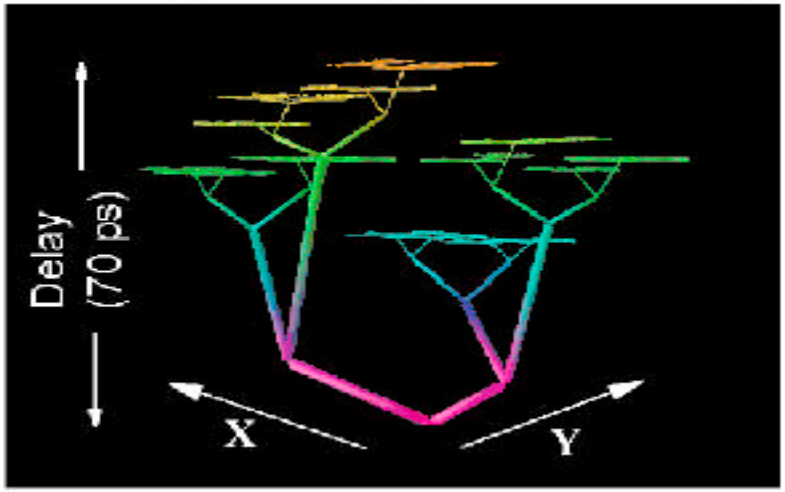

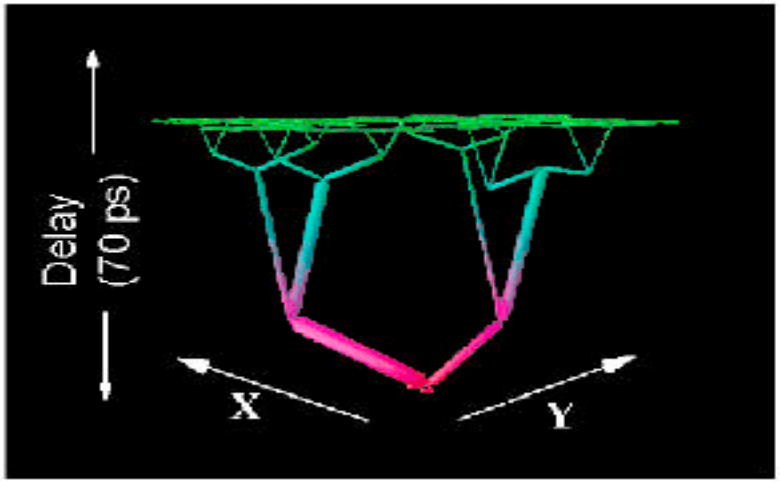

Clock circuits live in center column.

32 global clock wires go down the red column.

Any 10 may be sent to a clock region.

Also, 4 regional clocks (restricted functionality).

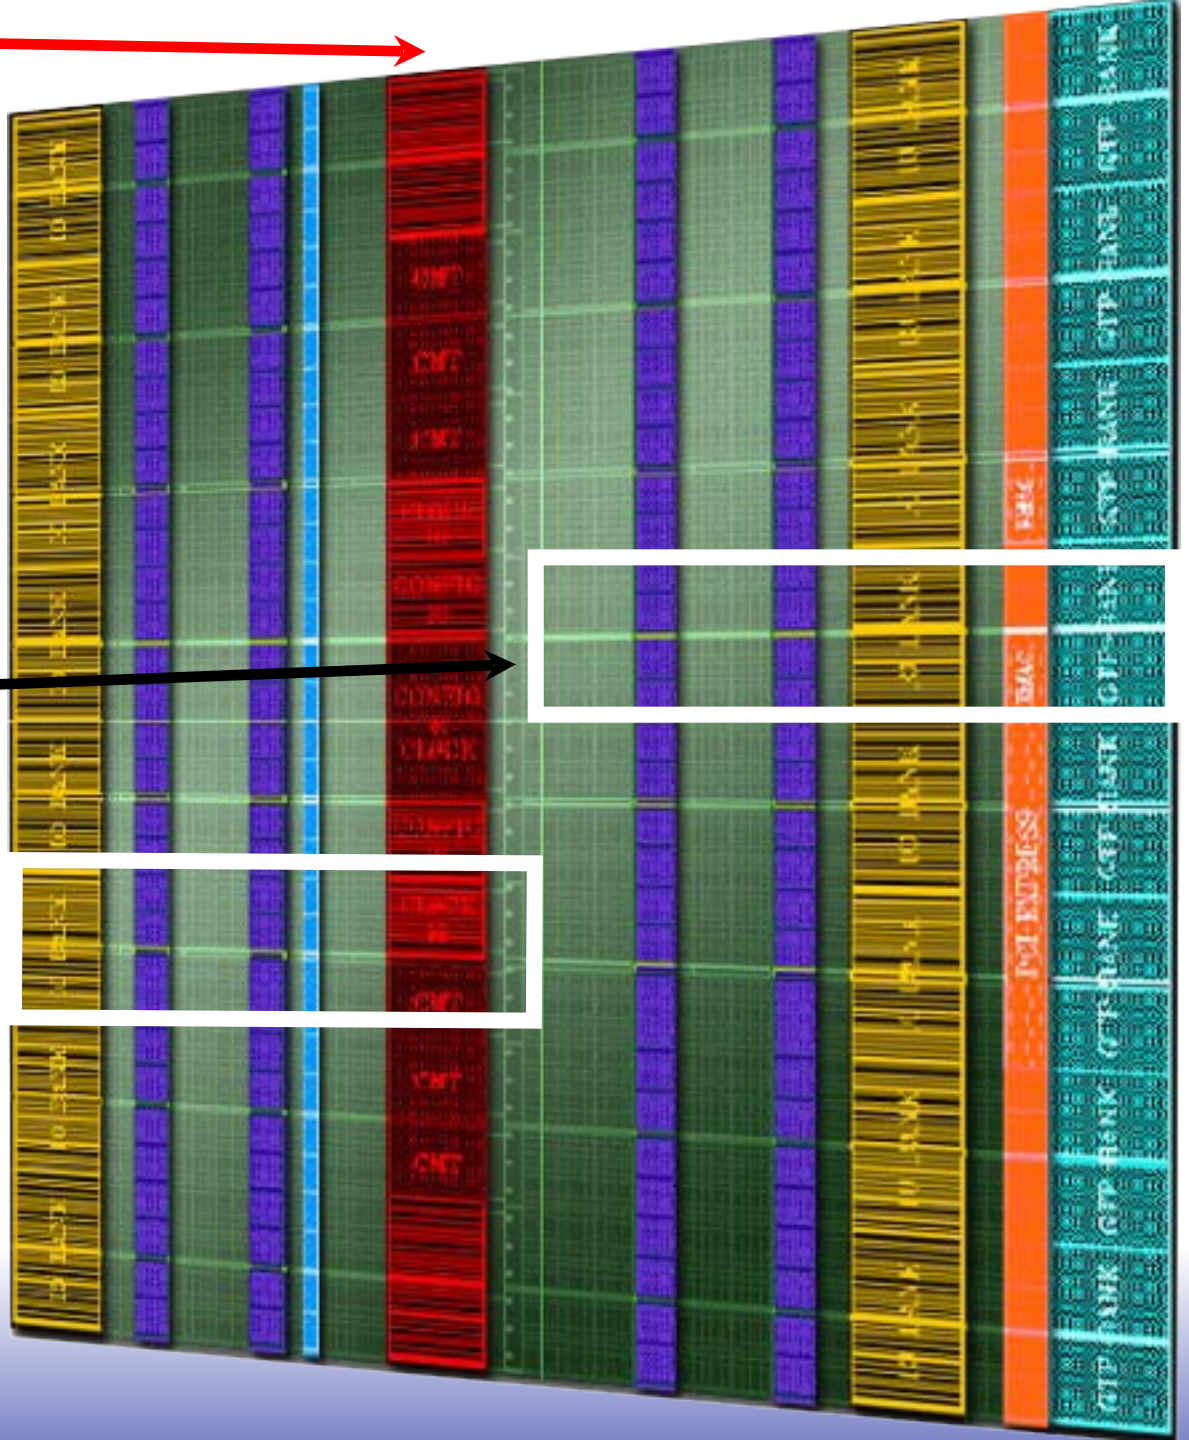

Clocks have dedicated wires (low skew)

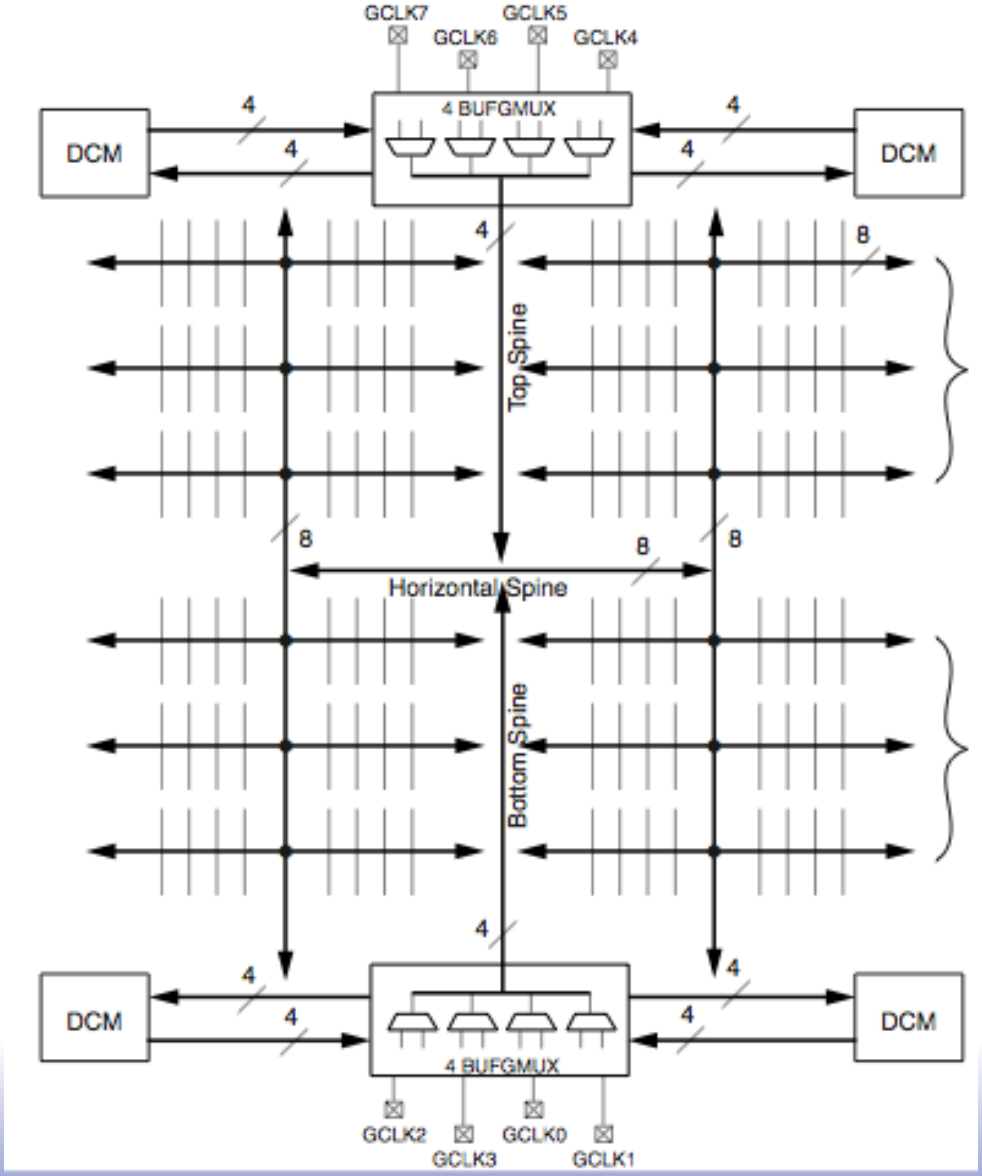

From: Xilinx Spartan 3 data sheet. Virtex is similar.

Low-skew Clocking in FPGAs

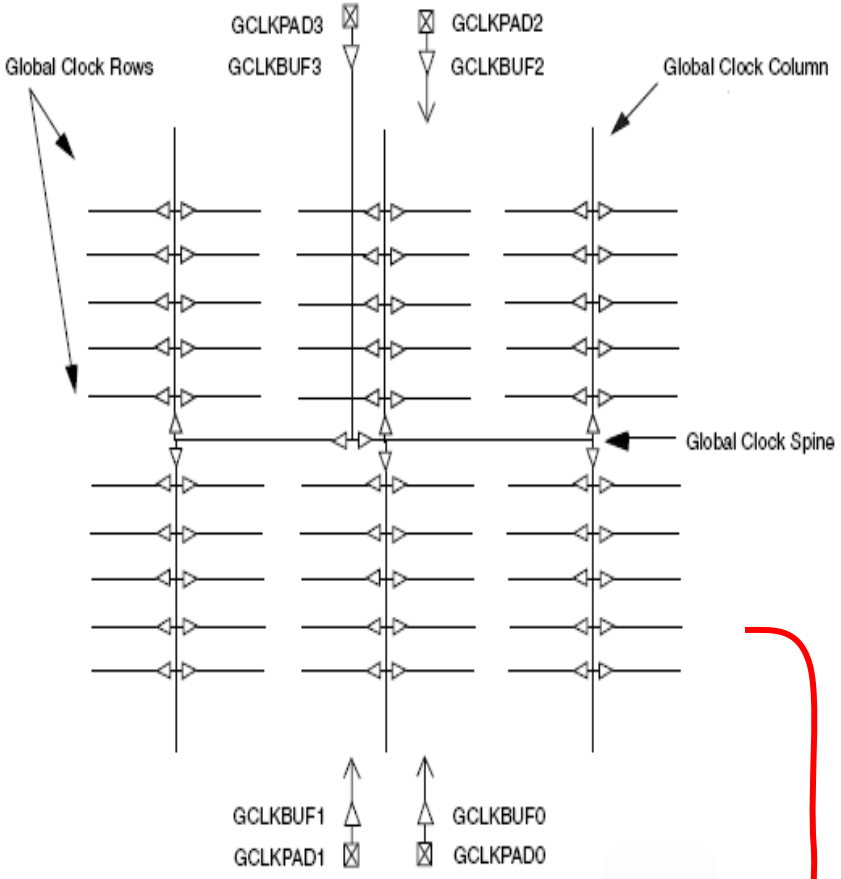

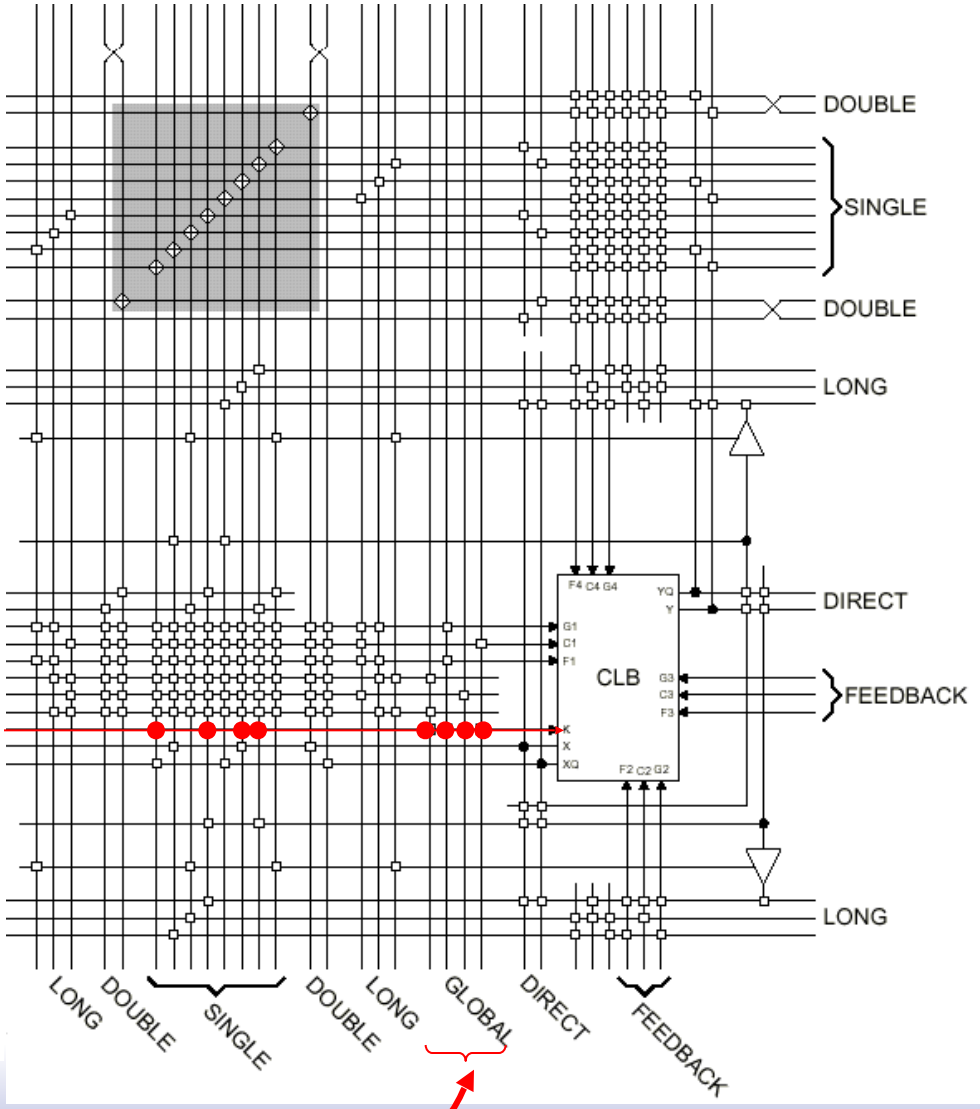

Figures from Xilinx App Notes

Die photo: Xilinx Virtex

Gold wires are the clock tree.

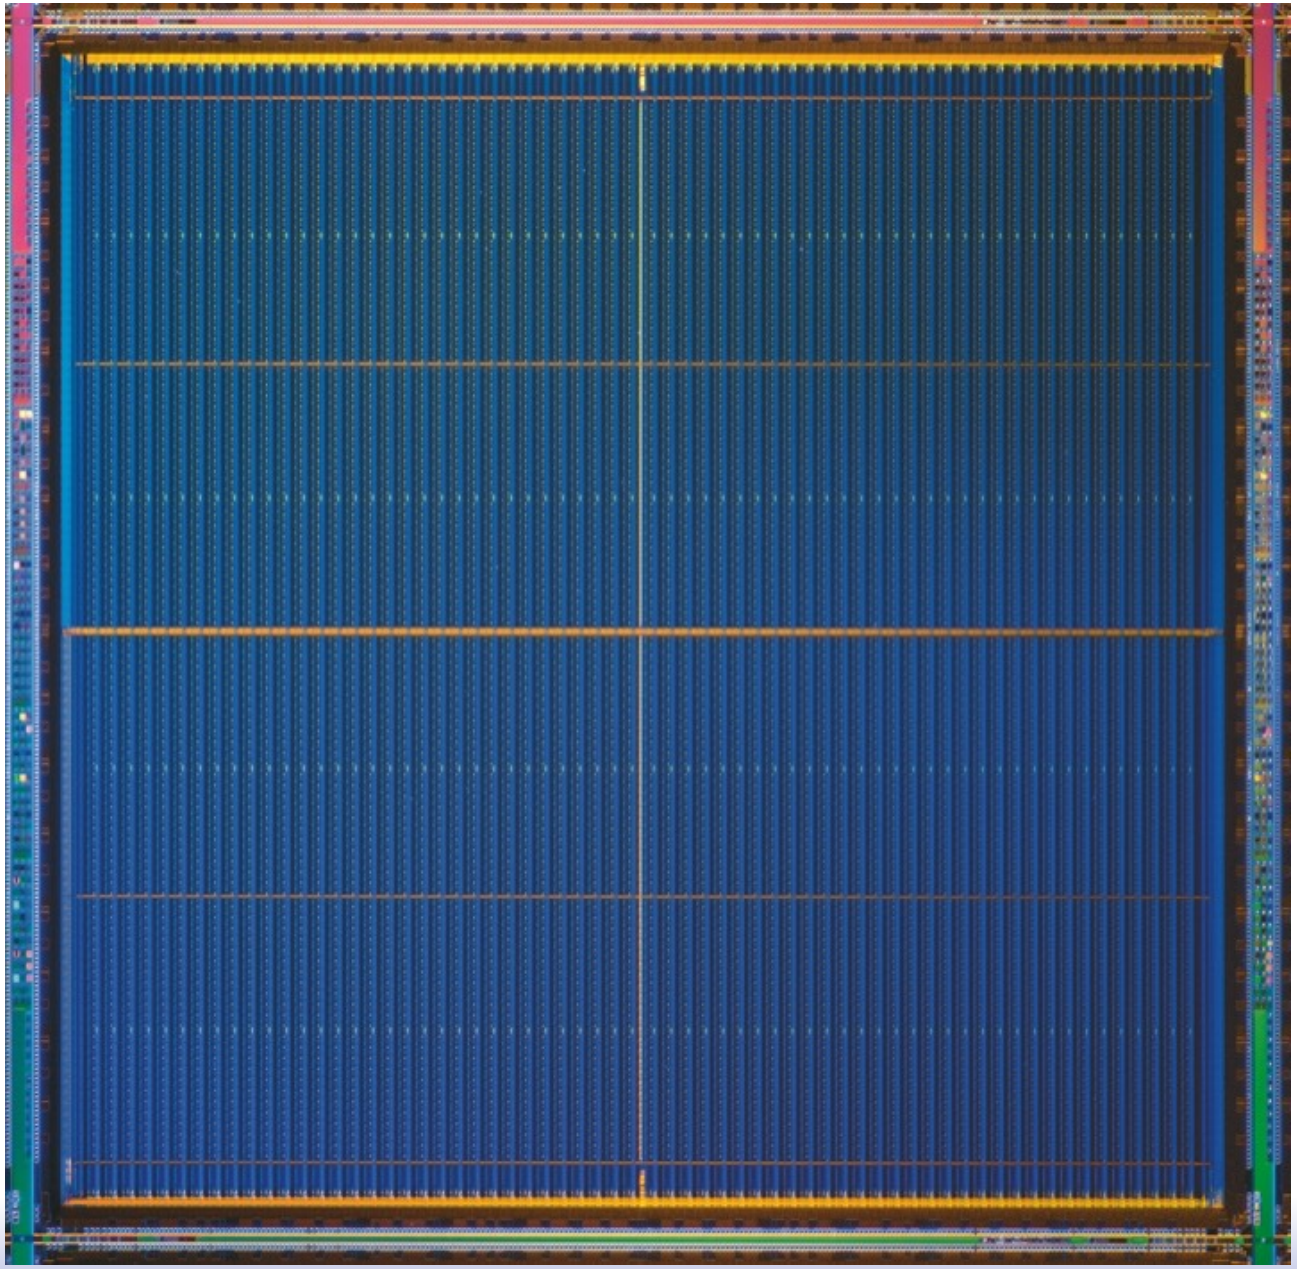

Chip Packaging

Chip Packaging

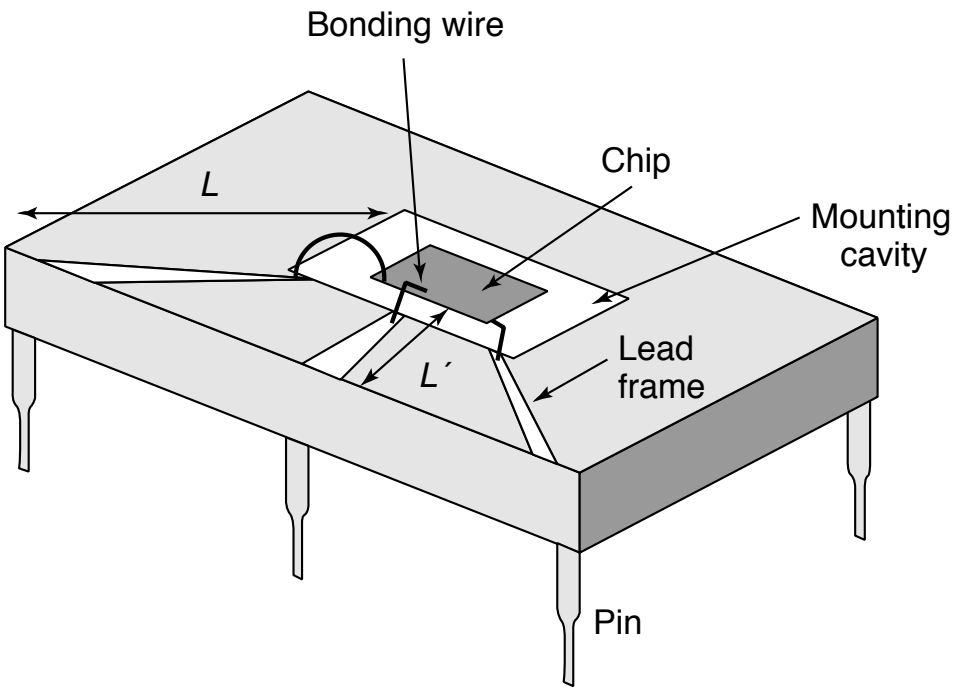

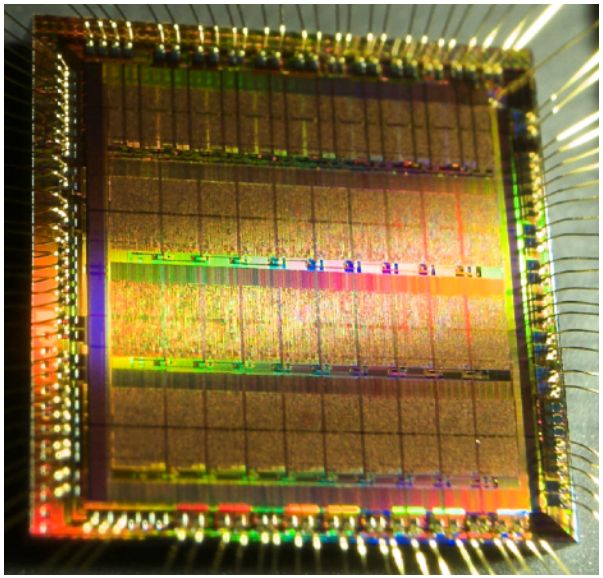

- Bond wires (~25μm) are used to connect the package to the chip
- Pads are arranged in a frame around the chip
- Pads are relatively large (~100μm in 0.25μm technology), with large pitch (100μm)
- 60μm x 80μm at 80μm pitch in 45nm
- Many chips are 'pad limited'

Pad Frame

Layout

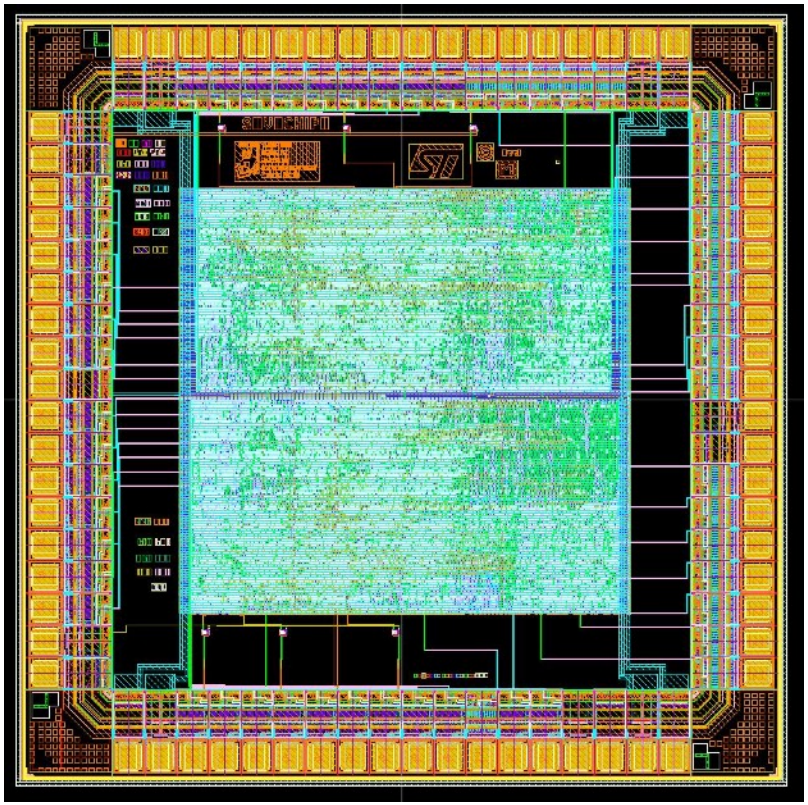

Die Photo

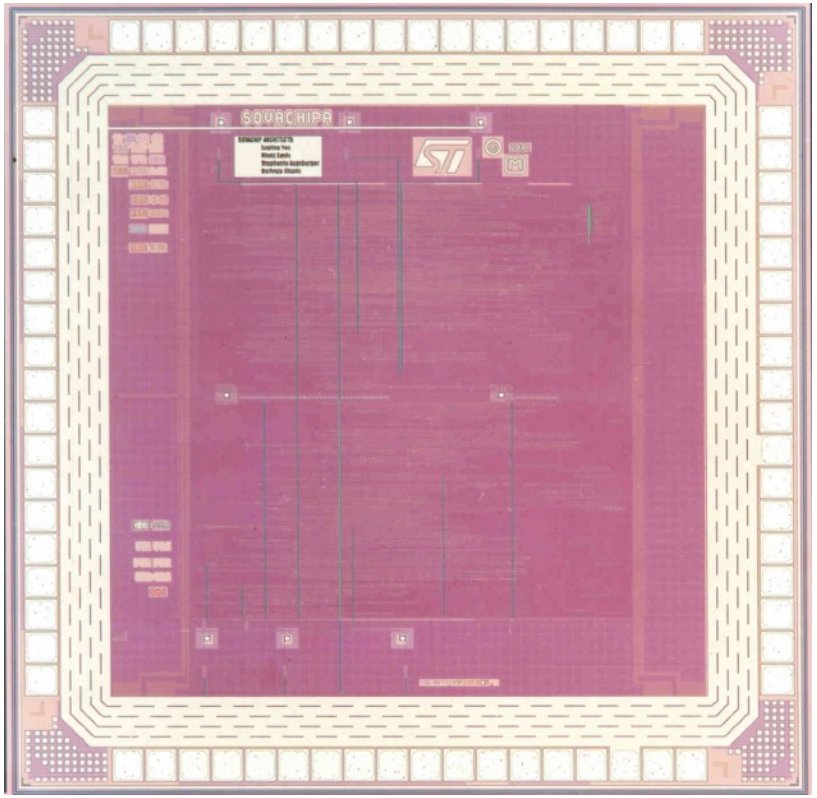

Bonding Pad Design

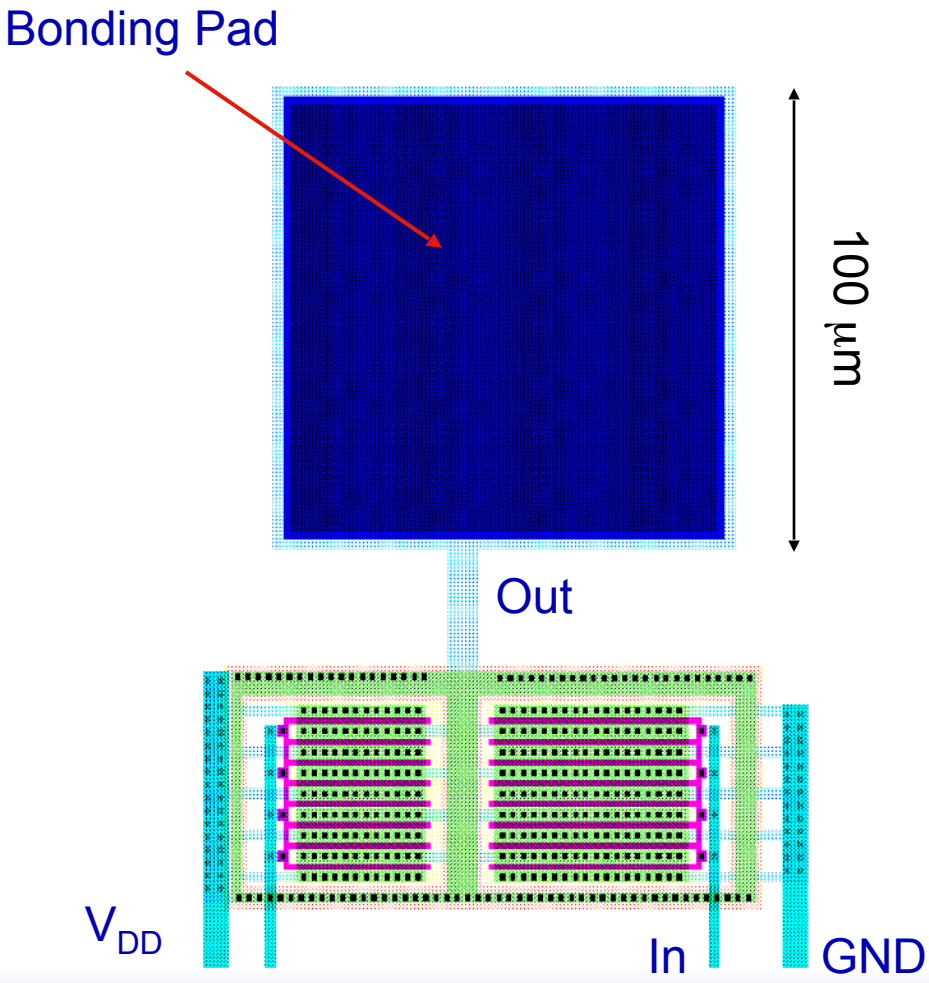

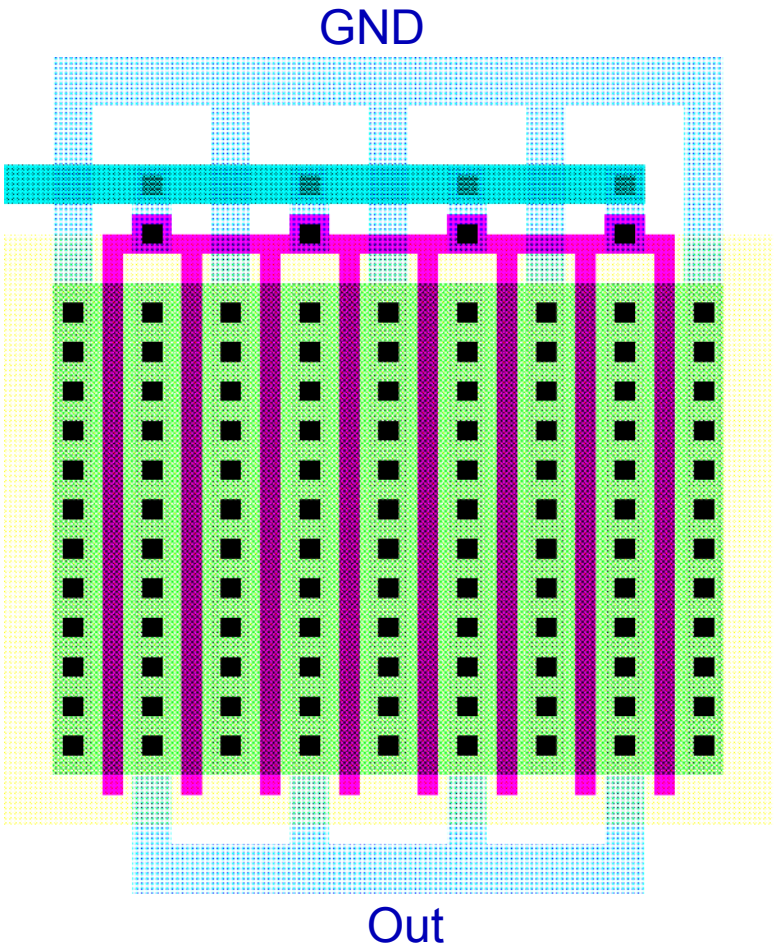

Chip Packaging

- 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

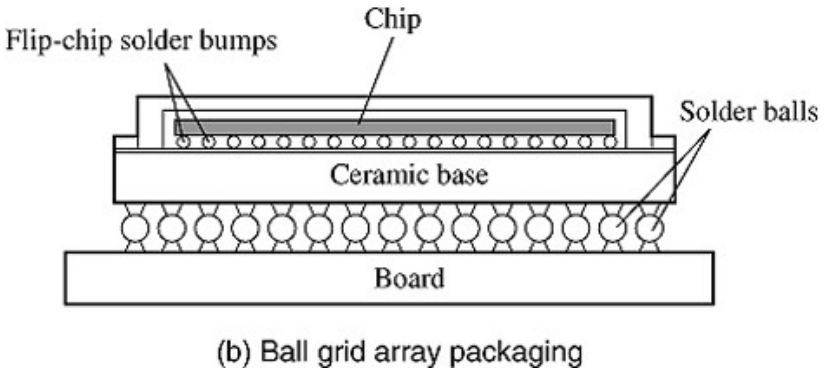

(b) Ball grid array packaging

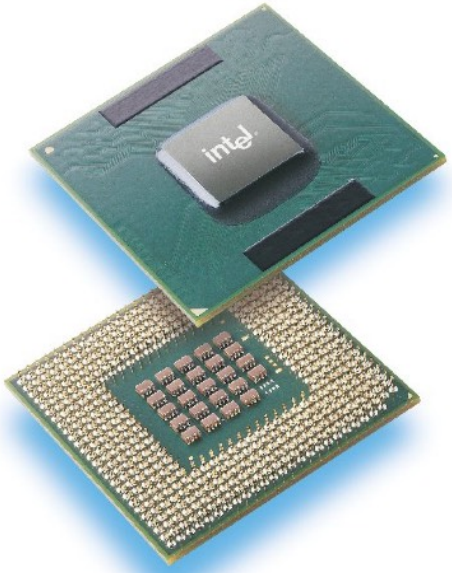

ESD Protection

- 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

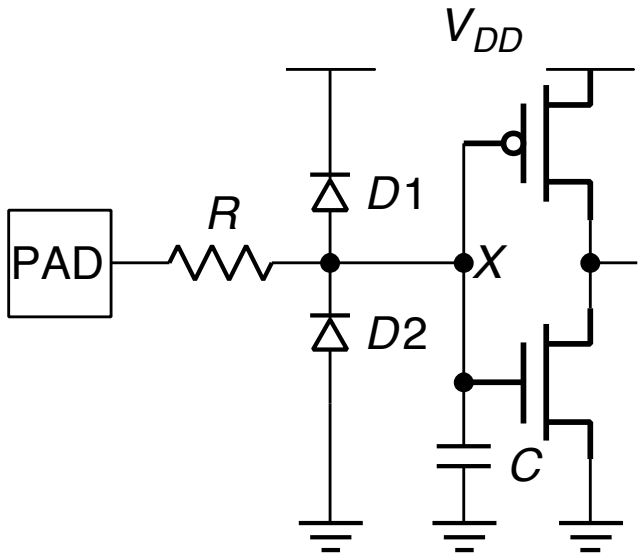

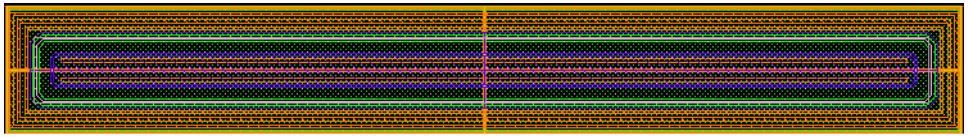

Diode

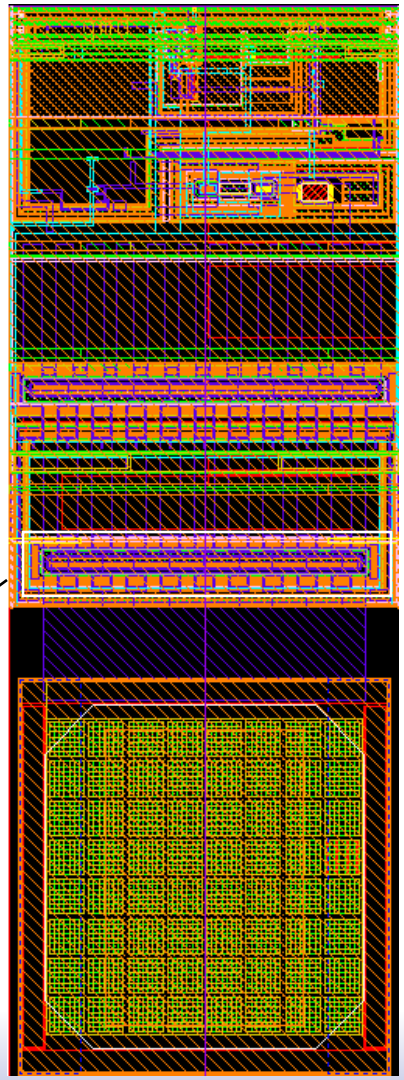

When Things Go Bad

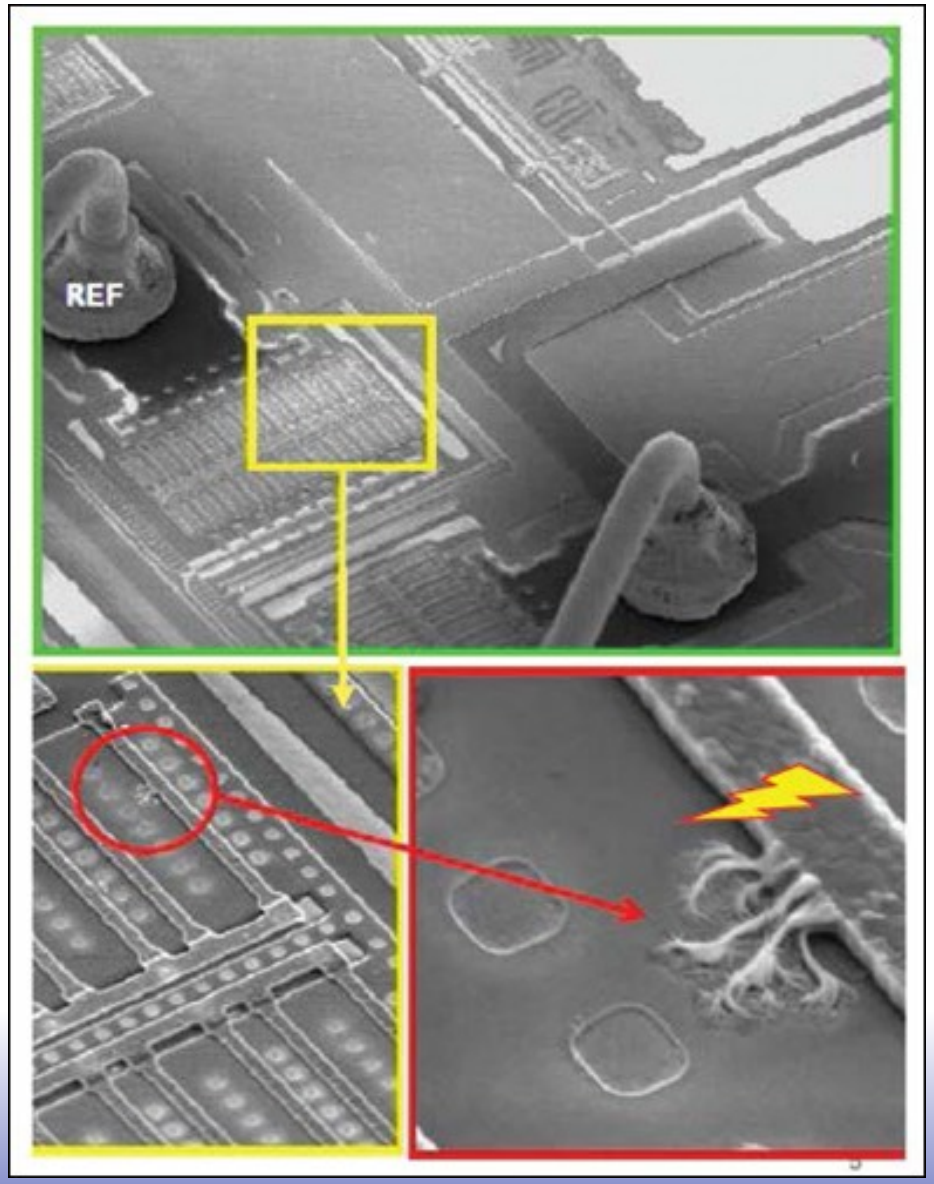

[Maxim]

Stray Paths

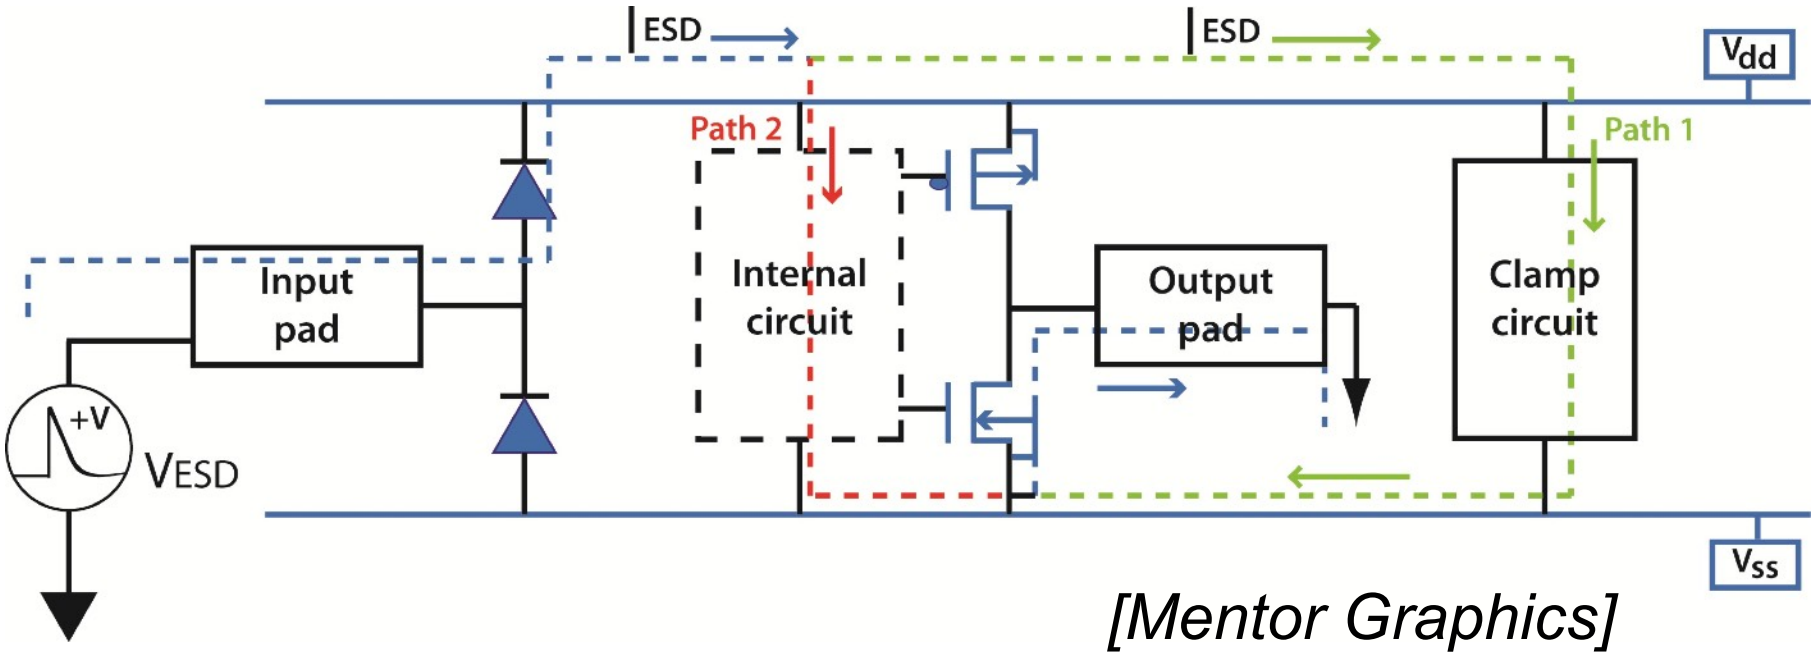

[Mentor Graphics]

- Clamp circuit sinks the ESD current
  - Typically a few stacked diodes (2-4)

Split Ground Rails

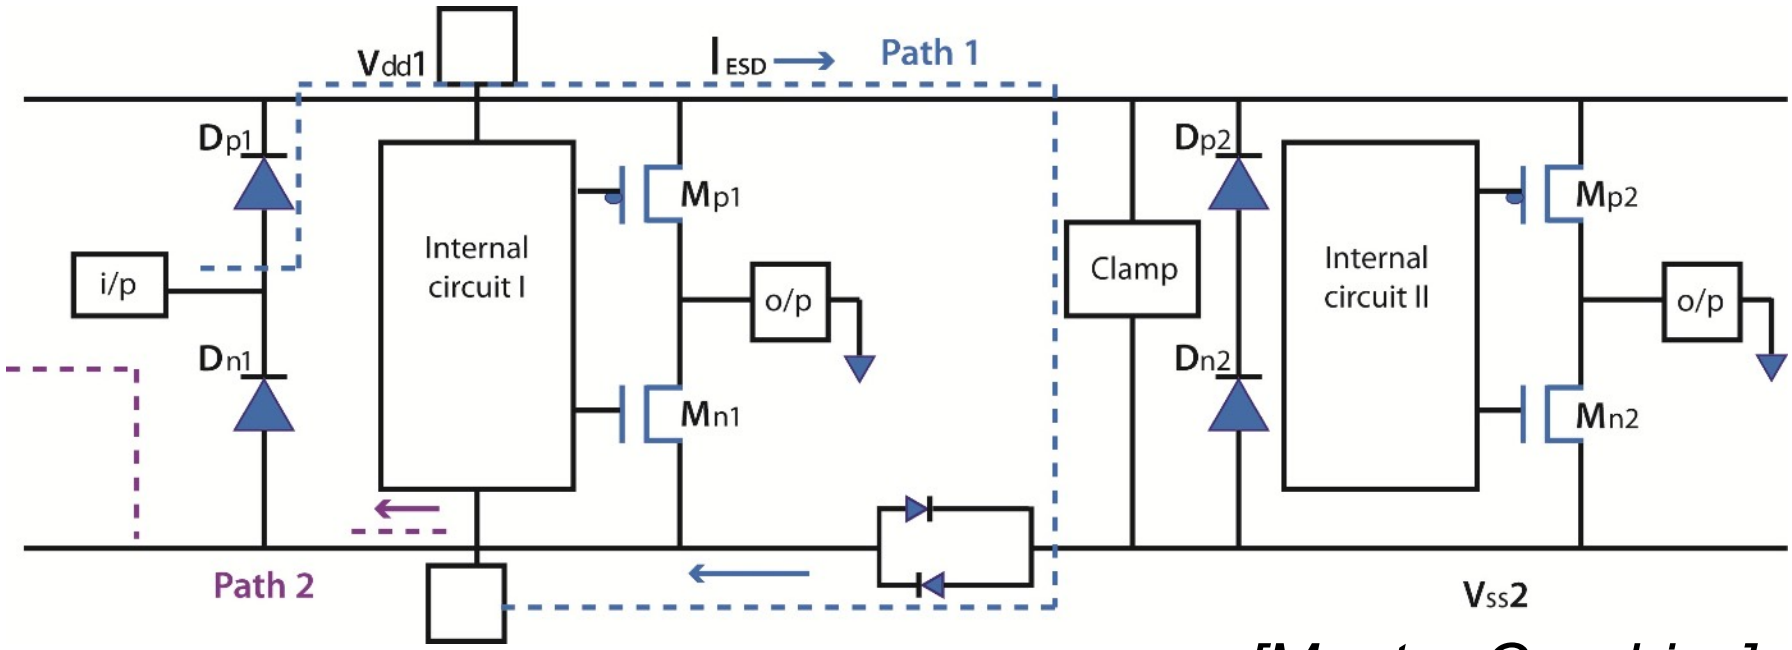

[Mentor Graphics]

- Bidirectional diode fixes the ESD path

Power Distribution

Power Supply Impedance

- No voltage source is ideal - ||Z|| > 0
- Two principal elements increase Z:
  - Resistance of supply lines (IR drop)
  - Inductance of supply lines (L·di/dt drop)

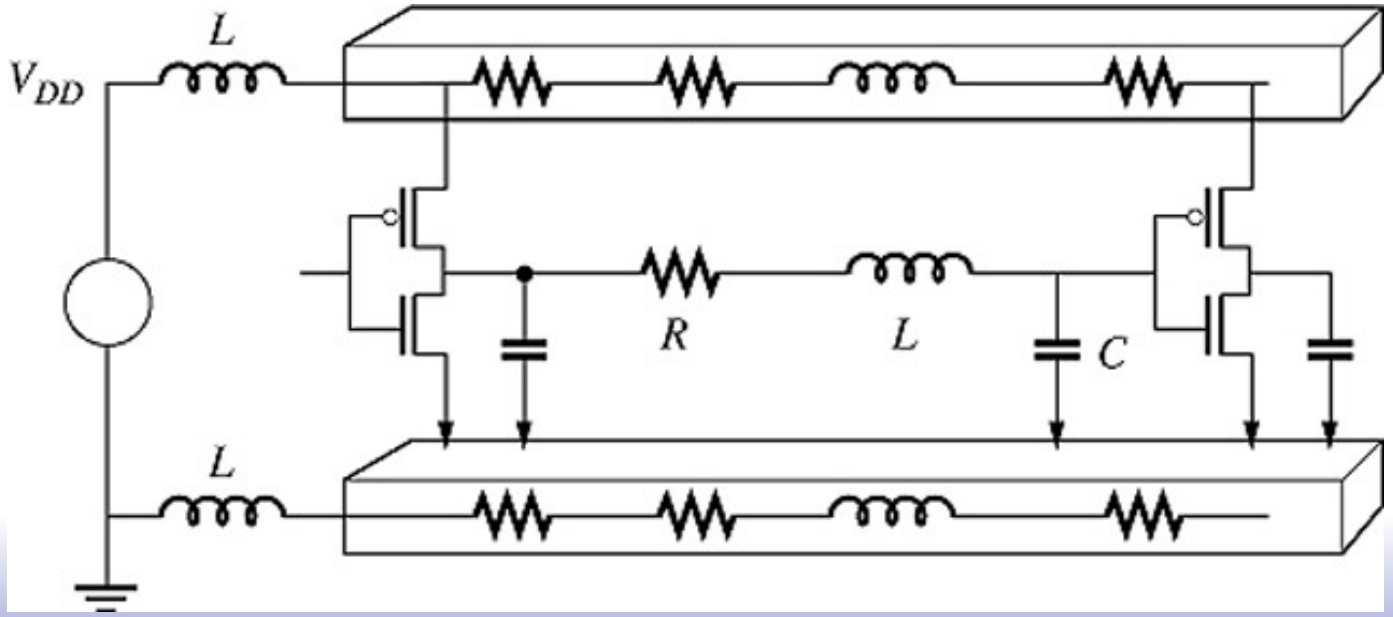

Scaling and Supply Impedance

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

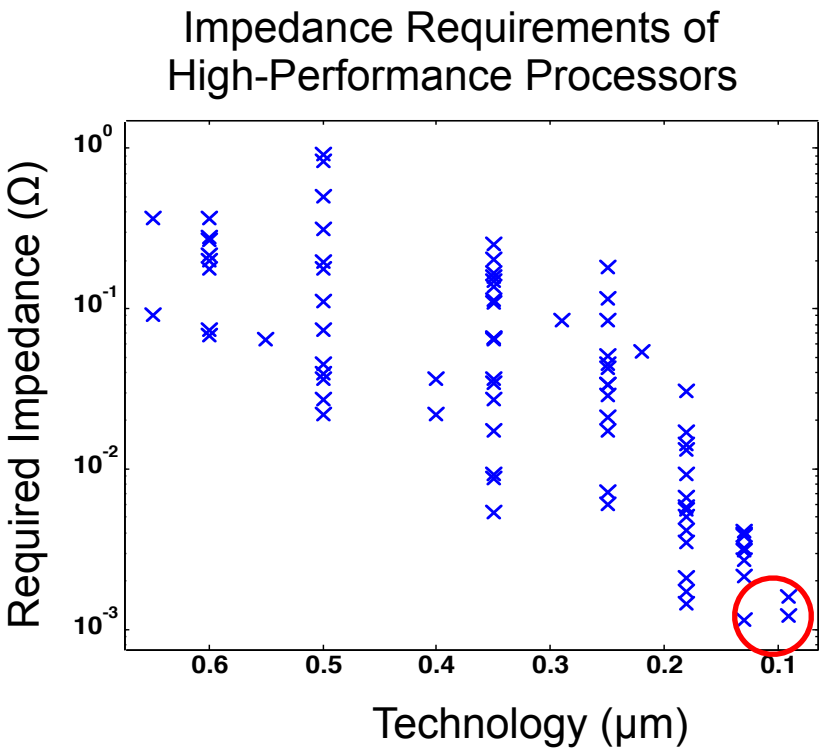

- In traditional scaling V_{dd} drops while power stays constant
- This forced drastic drop in supply impedance
  - $V_{dd} \downarrow$, $I_{dd} \uparrow \rightarrow \quad |Z_{required}| \downarrow\downarrow$
- Today's chips:
  - $|Z_{required}| \approx 1\ m\Omega$!
  - Vdd = 1V, P=100W => Idd=100A
  - For $\Delta V_{dd,max}$ = 100mV, $Z_{dd,max}$ = 100mV/100A = 1mΩ!

IR Drop Example

- Intel Pentium 4: ~103W at ~1.275V
  - I_{dd} = 81Amps
- For 10% IR drop, total distribution resistance must be less than 1.6mΩ
- On-chip wire R ≈ 20mΩ/sq. (thick metal)
  - Can't meet R requirement even with multiple, complete layers dedicated to power
  - Main motivation for flip-chip packaging

Power Delivery

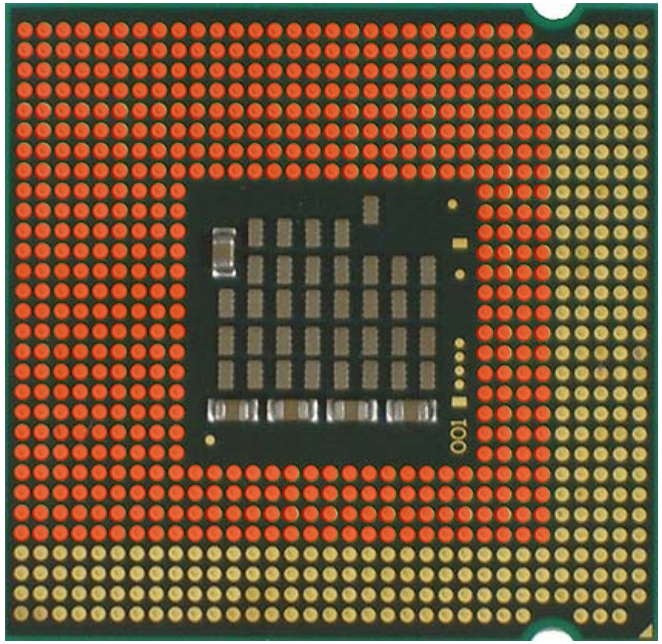

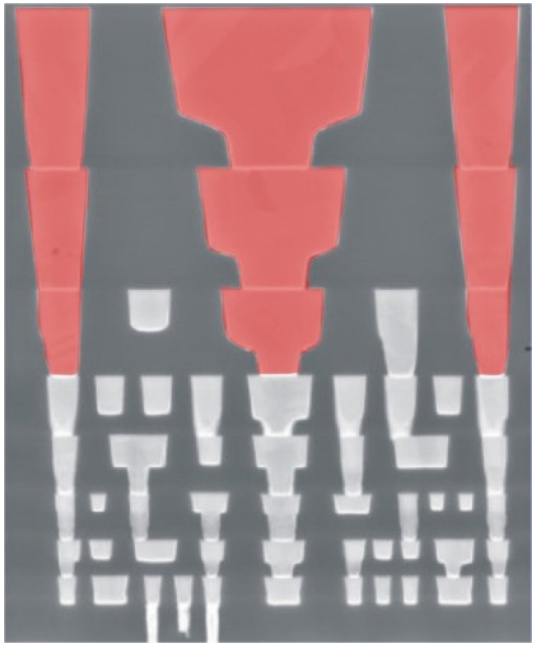

- Achieving such low impedance requires a lot of resources:
  - ~70% of package pins just for power
  - Top 2-3 (thick) metal layers

Not Just Impedance - Electromigration

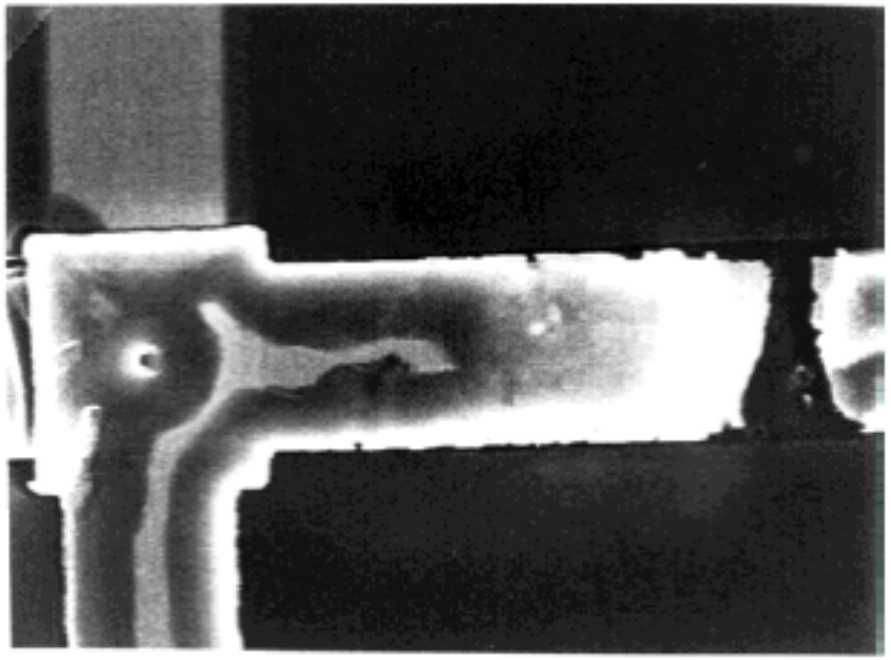

- On-chip wires: current limited to ~1mA/μm for 5-7 year lifetime

On-Chip Power Distribution

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

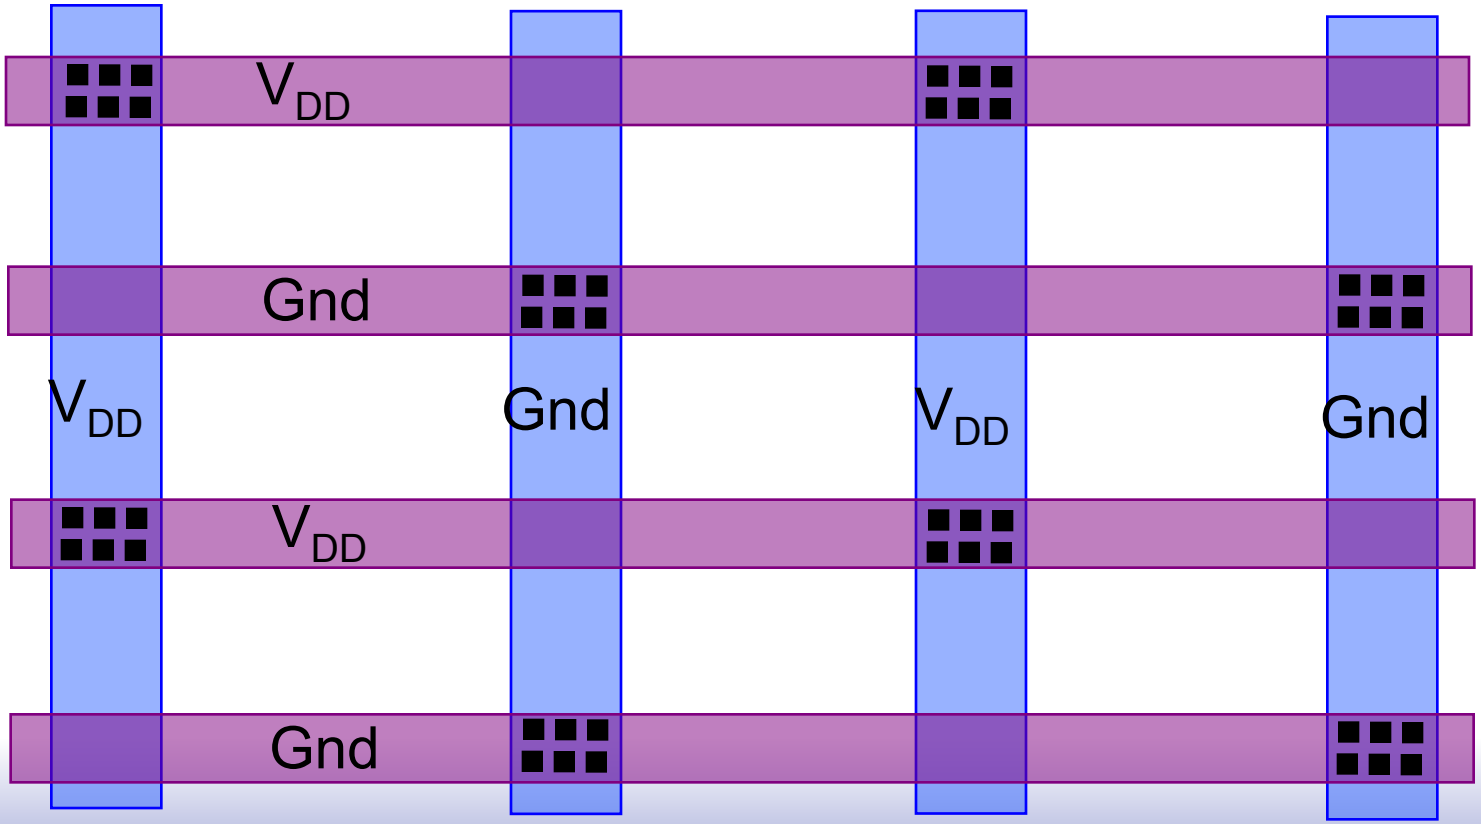

3 Metal Layer Approach (EV4)

3rd "coarse and thick" metal layer added to the technology for EV4 design

Power supplied from two sides of the die via 3rd metal layer

2nd metal layer used to form power grid

90% of 3rd metal layer used for power/clock routing

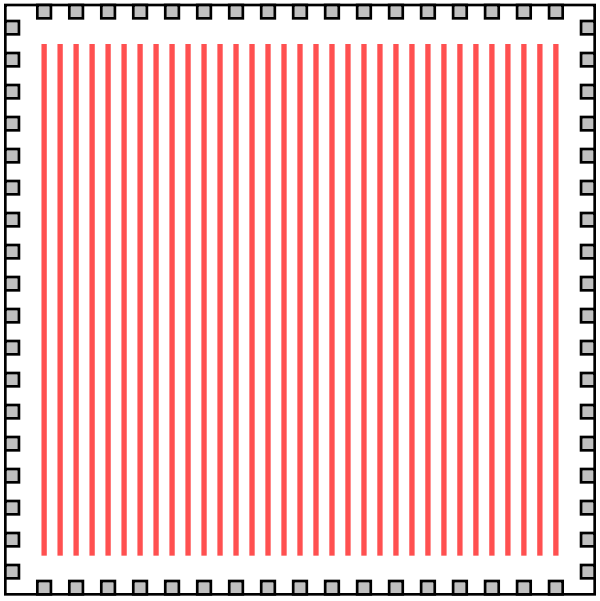

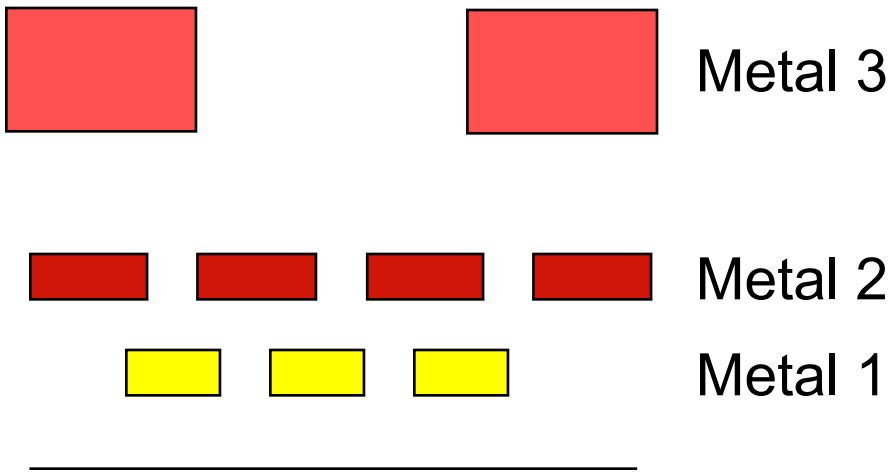

Courtesy Compaq

4 Metal Layers Approach (EV5)

4th “coarse and thick” metal layer added to the technology for EV5 design

Power supplied from four sides of the die

Grid strapping done all in coarse metal

90% of 3rd and 4th metals used for power/clock routing

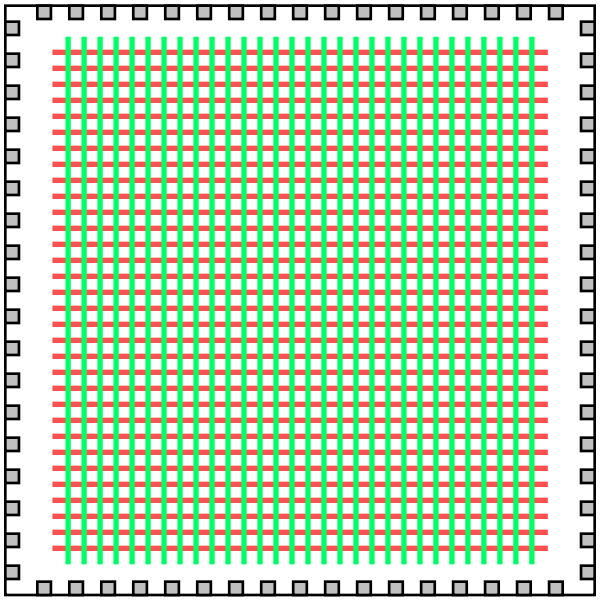

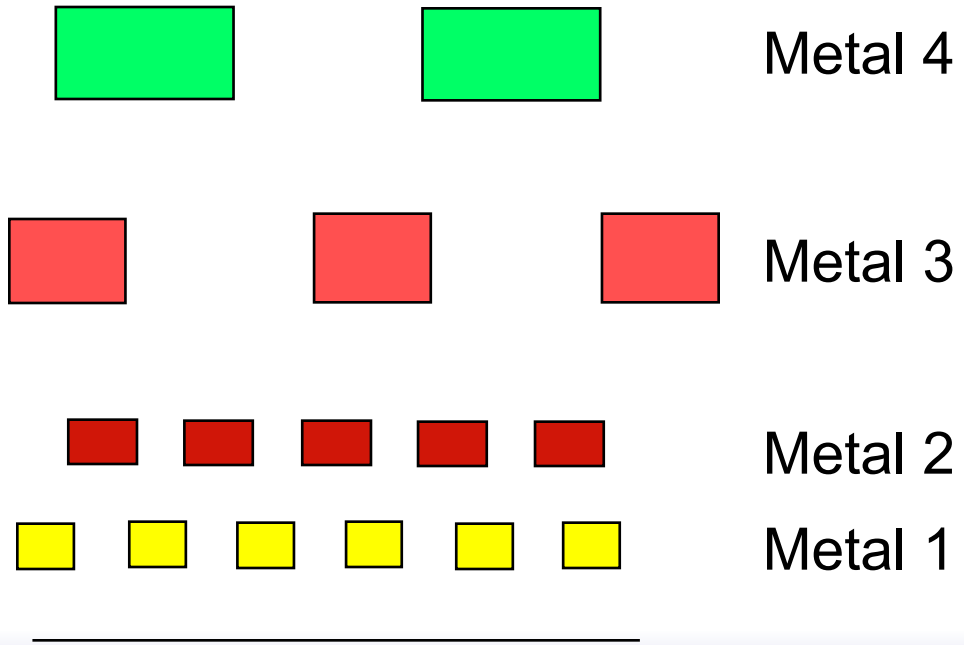

Courtesy Compaq

6 Metal Layer Approach – EV6

2 reference plane metal layers added to the technology for EV6 design

Solid planes dedicated to Vdd/Vss

Lowers on-chip inductance

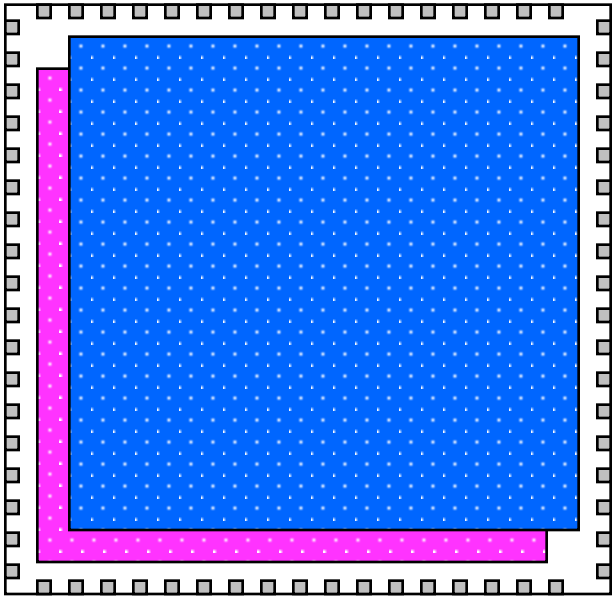

Courtesy Compaq

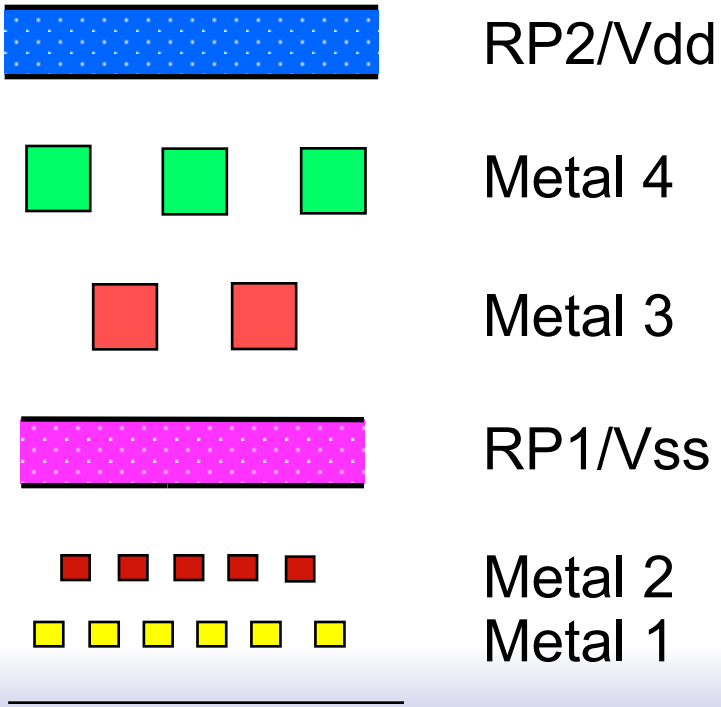

Decoupling Capacitors

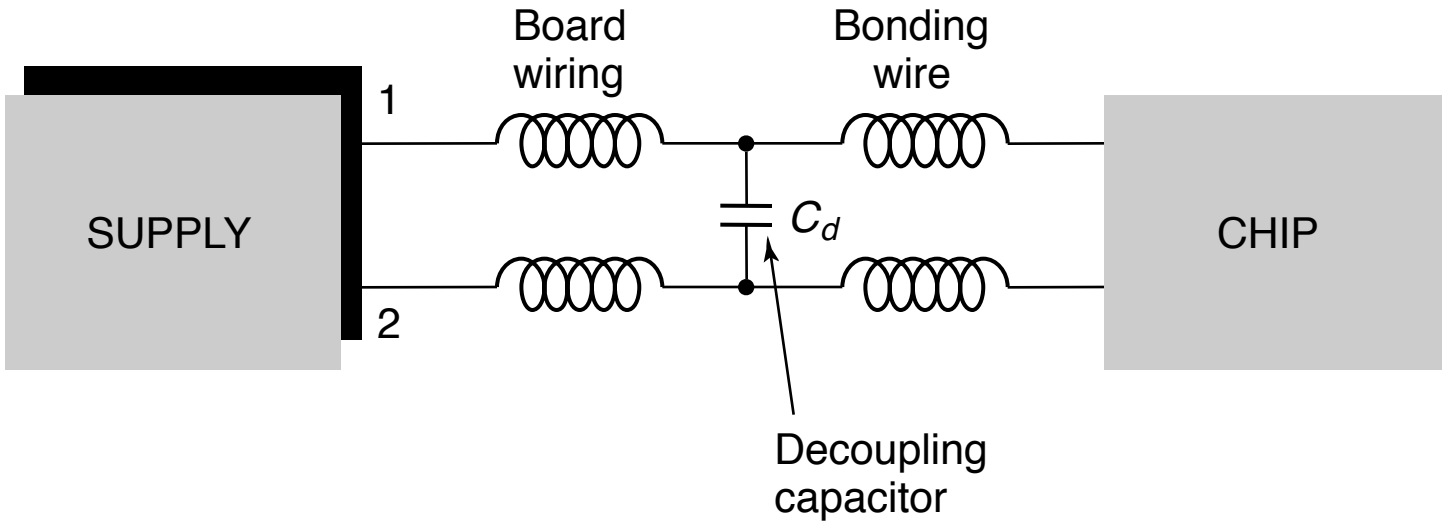

Decoupling capacitors are added:

- On the board (right under the supply pins)
- On the chip (under the supply straps, near large buffers)

Decoupling Capacitors

- Under the die

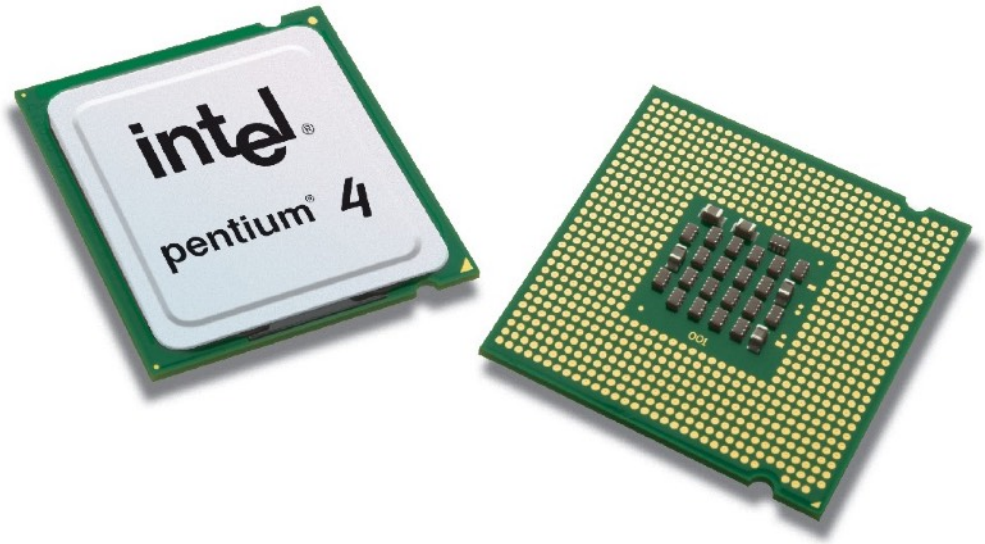

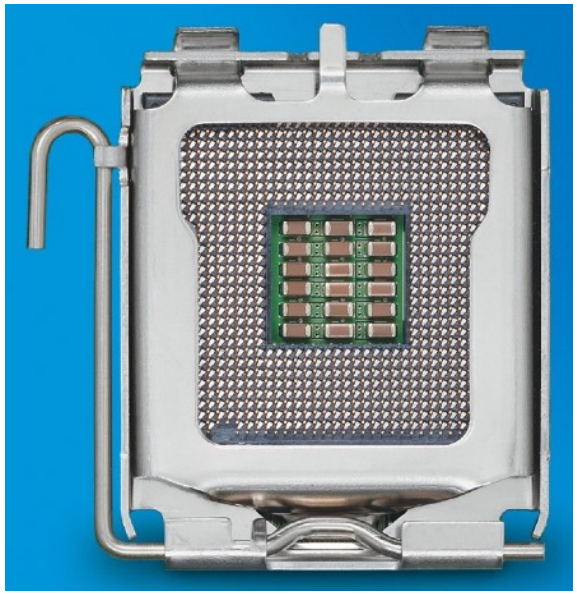

Pin Inductance Example

- Processor transient current is 100A in 20ps from 1V supply
- C4 bump inductance is 25pH
- How many C4 bumps do we need to get supply noise spike of less than 10%?

- With wirebond inductance of 1nH (1nH/mm) how many wirebonds are needed?

Modern Concepts: Integrated Regulators

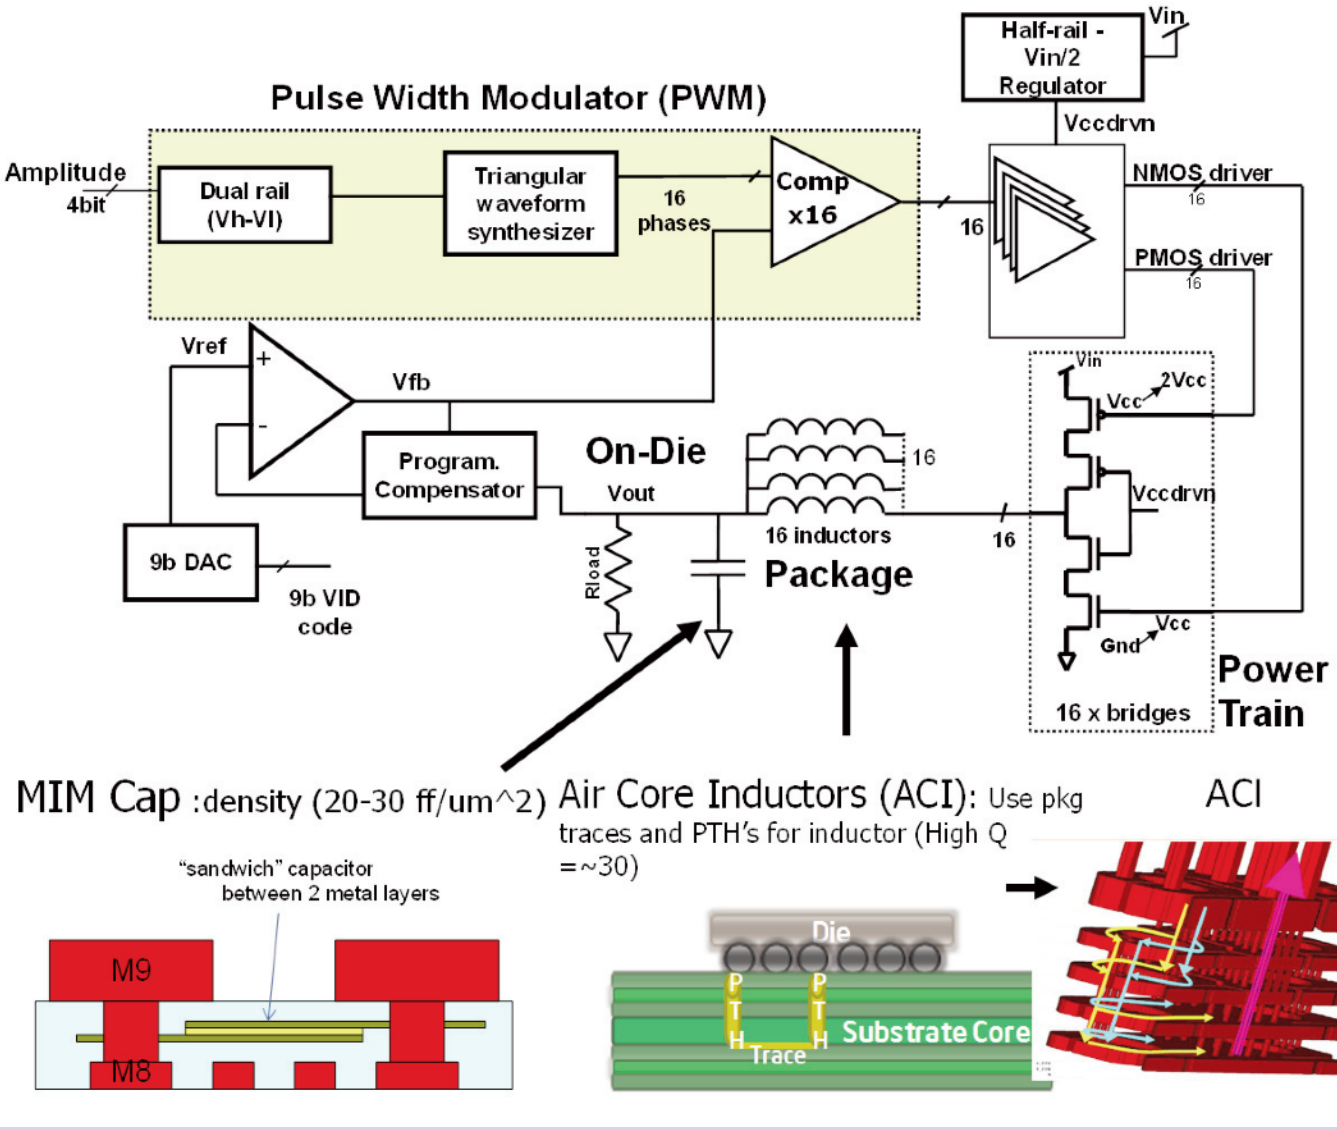

- 1.7V down to 1V, 0.1-16A *[Haswell CPU]*