EECS 427
Lecture 21: Clock distribution & robustness

Reading: 10.3.3

Last Time

- Advanced low-power techniques, especially static power reduction
  - Dual-Vth, can save ~80%, ΔVth ~ 100mV
  - Stack effect, understand this
  - Sleep transistors; how they operate
  - Bias transistor body to alter Vth through body effect
- Main idea: gates with timing slack can be slowed down in various ways that help reduce power

Lecture Overview

- Clock distribution metrics
  - Area, power, skew
- Clock network distribution types
  - Tree and grid, hybrid
- Recall clock skew and jitter from previous lectures

Clocks: Power-Hungry

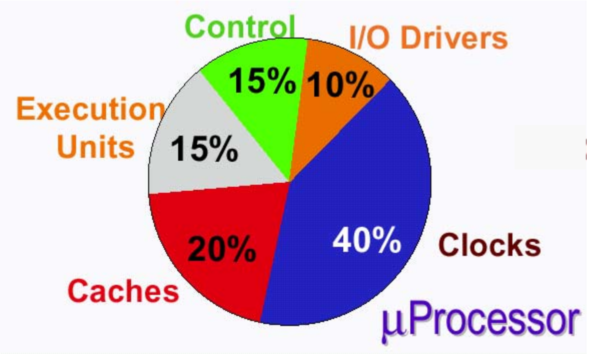

$$P = \alpha C V_{dd}^2 f$$

Not only is the clock capacitance large, it switches every cycle!

Clock Distribution Metric: Area

- Clock networks consume silicon area (clock drivers, PLL, etc.) and routing area
- Routing area is most vital
- Top-level metals are used to reduce RC delays
  - These levels are precious resources (unscaled)
  - Power routing, clock routing, key global signals
- By minimizing area used, we also reduce wiring capacitance & power
- Typical #'s: Intel Itanium – 4% of M4/5 used in clock routing

---

Slew Rates

- To maintain signal integrity and latch performance, minimum slew rates are required
- Too slow – clock is more susceptible to noise, latches are slowed down, eats into timing budget
- Too fast – burning too much power, overdesigned network, enhanced ground bounce
- Rule-of-thumb: T_{rise} and T_{fall} of clock are each between 10-20% of clock period (10% - aggressive target)
  - 1 GHz clock; $T_{rise} = T_{fall} = 100\text{-}200\text{ps}$

Technology Trends: Power

- Heavily pipelined designs ⇒ more latches, more capacitive load for clock
- Larger chips ⇒ more wirelength needed to cover the entire die
- Complexity ⇒ more functionality and devices means more clocked elements
- Dynamic logic ⇒ more clocked elements

The Grid System

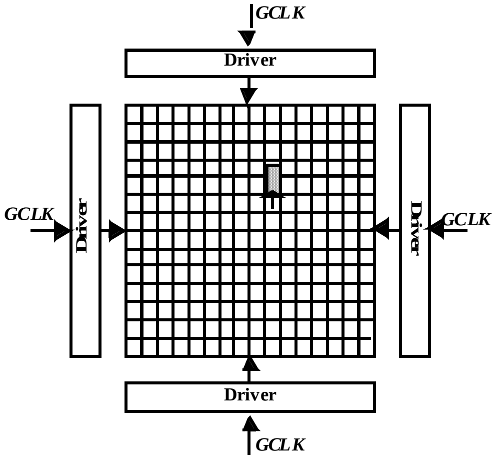

- No matching
- Large power (huge drivers)

Network Types: Grid

- Gridded clock distribution was common on earlier DEC Alpha microprocessors
- Advantages:
  - Skew determined by grid density and not overly sensitive to load position
  - Clock signals are available everywhere
  - Tolerant to process variations
  - Usually yields extremely low skew values

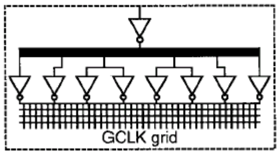

Pre-drivers

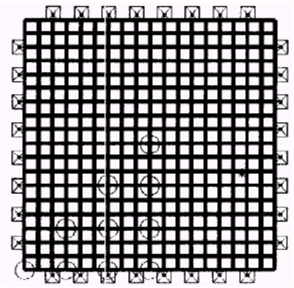

Global grid

Grid Disadvantages

- Huge amounts of wiring & power
  - Wire cap large
  - Strong drivers needed – pre-driver cap large
  - Routing area large
- To minimize all these penalties, make grid pitch coarser
  - Skew gets worse
  - Losing the main advantage
- Don’t overdesign – let the skew be as large as tolerable
- Grids aren’t feasible for most designs due to power

Network Types: Tree

- Original H-tree (Bakoglu)
    - One large central driver
    - Recursive H-style structure to match wirelengths
    - Halve wire width at branching points to reduce reflections

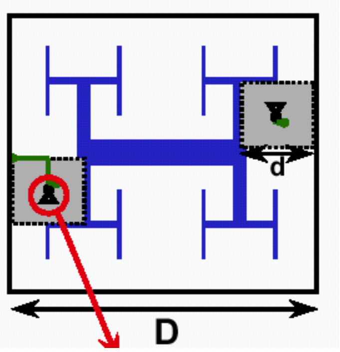

H-Tree Problems

- Drawback to original tree concept
    - slew degradation along long RC paths
    - unrealistically large central driver
        - Clock drivers can create large temperature gradients (ex. Alpha 21064 ~30° C)
    - non-uniform load distribution
- Inherently non-scalable (wire resistance skyrockets)
- Solution to some problems
    - Introduce intermediate buffers along the way
    - Specifically at branching points

Buffered H-tree

- Advantages
  - Ideally zero-skew
  - Can be low power (depending on skew requirements)
  - Low area (silicon and wiring)
  - CAD tool friendly (regular)
- Disadvantages
  - Sensitive to process variations
  - Local clocking loads are inherently non-uniform

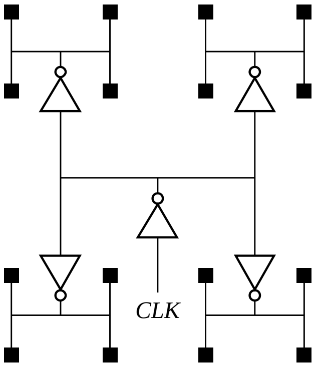

Realistic H-tree

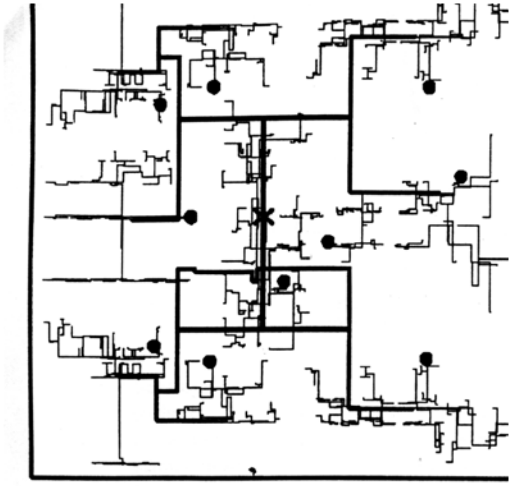

[Restle98]

Balancing a Tree

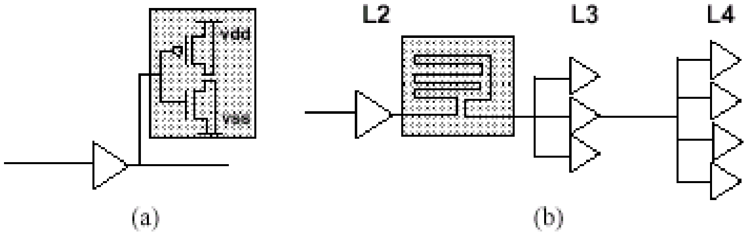

Con: Routing area often more valuable than silicon

Some techniques:

a) Introduce dummy loads

b) Snaking of wirelength to match delays

Network of choice in high-performance

- Globally – Tree
- Why?
- Power requirements are reduced compared to global grid
  - Smaller routing requirements, frees up global tracks
- Trees are easily balanced at the *global* level
  - Keeps global skew low (with minimal process variation)

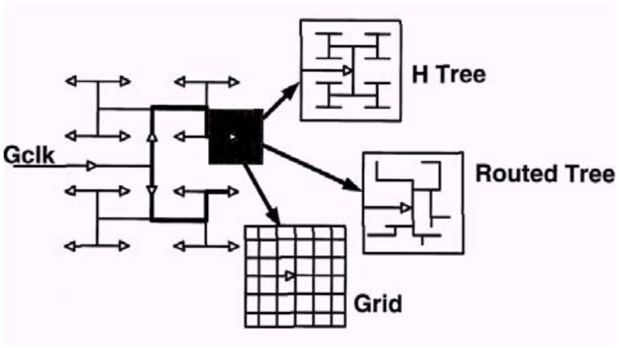

Network of Choice

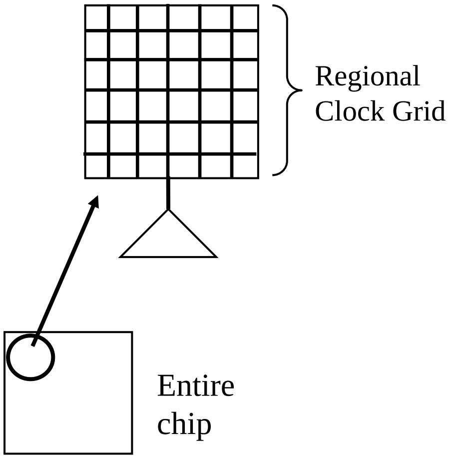

- Locally – Grid
- Why?
- Smaller grid distribution area allows for coarser grid pitch
  - Lower power in interconnect
  - Lower power in pre-drivers
  - Routing area reduced
- Local skew is kept very small
- Easy access to clock by simply tapping off grid

Example: DEC Alpha 21164

Clock Frequency: 300 MHz - 9.3 Million Transistors

Total Clock Load: 3.75 nF

Power in Clock Distribution network : 20 W (out of 50)

Uses Two Level Clock Distribution:

- **Single 6-stage driver at center of chip**
- **Secondary buffers drive left and right side clock grid in Metal3 and Metal4**

Total driver size: 58 cm!

21164 Clocking

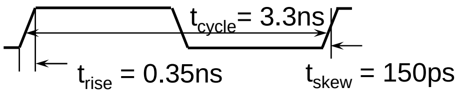

Clock waveform

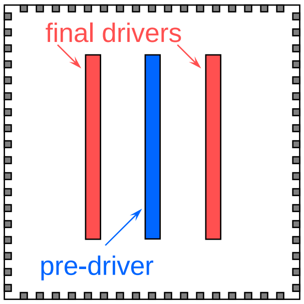

Location of clock driver on die

- 2 phase clock, distributed globally
- 2 distributed driver channels
  - Reduced RC delay/skew
  - Improved thermal distribution
  - 3.75nF clock load
  - 58 cm final driver width
- Local inverters for latching
- Conditional clocks in caches to reduce power
- More complex race checking
- Device variation

Clock Skew in Alpha Processor

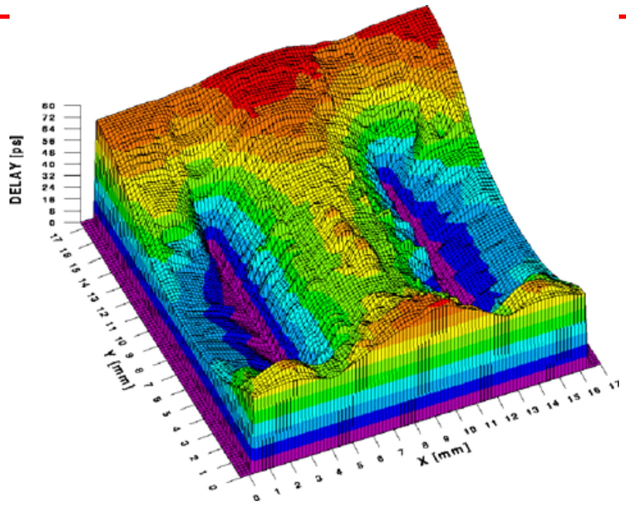

EV6 (Alpha 21264) Clocking
600 MHz – 0.35 micron CMOS

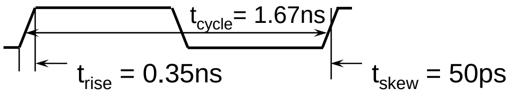

Global clock waveform

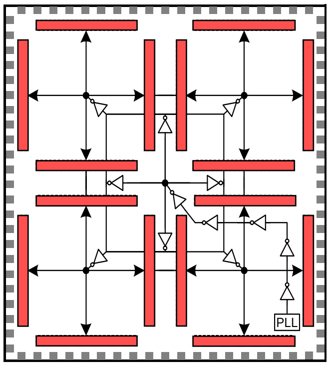

- 2.8 nF clock load
  - 40 cm final driver width
- Local clocks can be gated to save power
- Reduced load/skew
- Reduced thermal issues

EV6 Clock Results

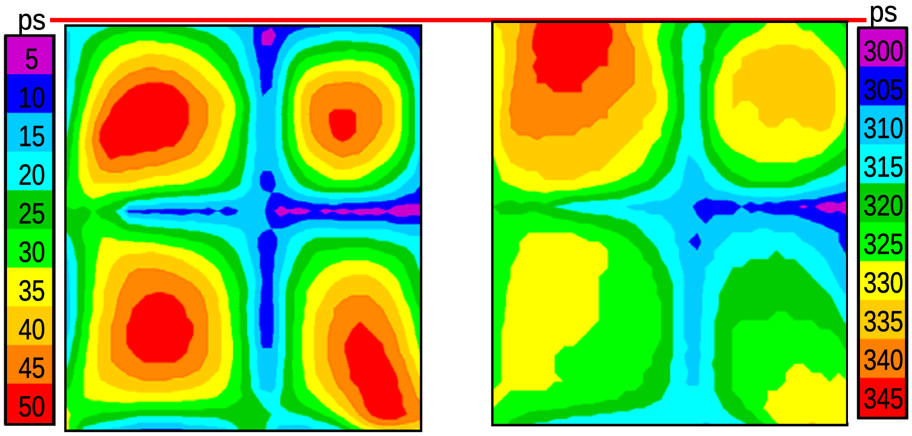

GCLK Skew
(at Vdd/2 Crossings)

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

Conclusions

- Getting the clock everywhere on a die at the exact same time is difficult
  - Requires a lot of power to reduce skew (big drivers, wide wires, etc.)
- Balanced H-trees are in common use
  - Design automation tools exist to synthesize these trees
- Clocks must be robust to variations/noise, have sharp slew rates, not create too much heat, plus other constraints