EECS 427
Lecture 22: Clock and Power Grid
Readings: 10.3.3, WH 12.3, CBF Ch.24

Reminders

- Project completion by Monday Dec. 14
  - Presentation in class: 1:40-3:30 (~15 minutes per group)
  - Project demo in CSE 1620: 3:30-5:30 (~15 minutes per group)
  - Voluntary individual contribution survey (Dec. 10 – 13)
  - HW5 (final report) due 7 pm
  - CAD9 due 7 pm
- Extra office hours
- Verify your grades on Monday

Clocking

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

Clocks: Power-Hungry

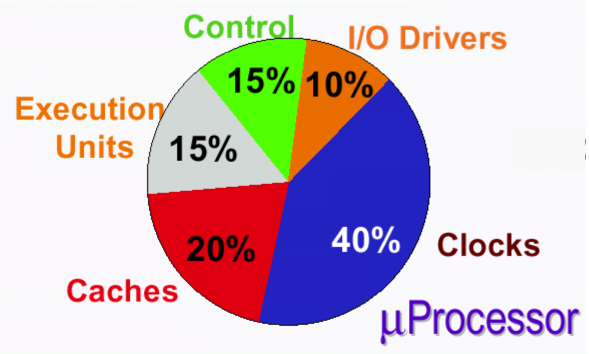

$$P = \alpha C V_{dd}^2 f$$

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; $T_{rise} = T_{fall} = 100\text{-}200$ps

The Grid System

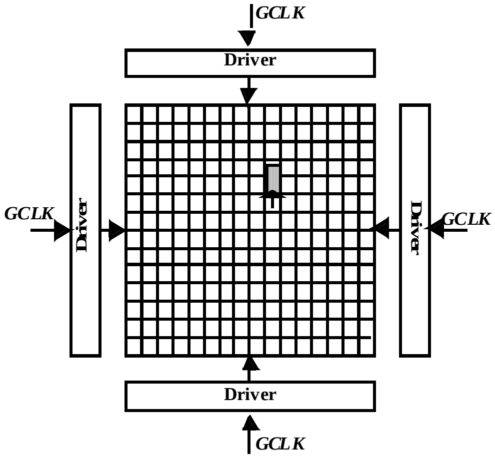

- No matching
- Large power (huge drivers)

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

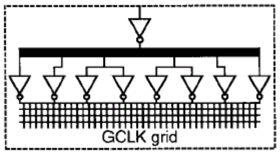

Pre-drivers

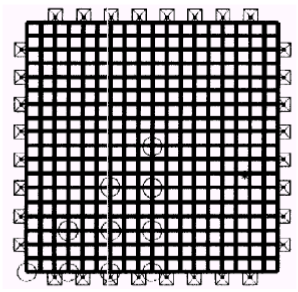

Global grid

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

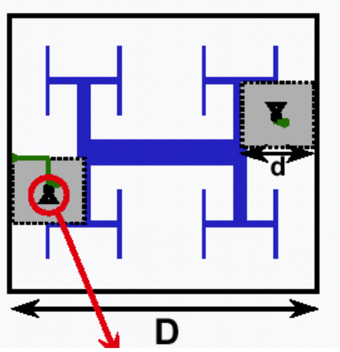

H-Tree Problems

- 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

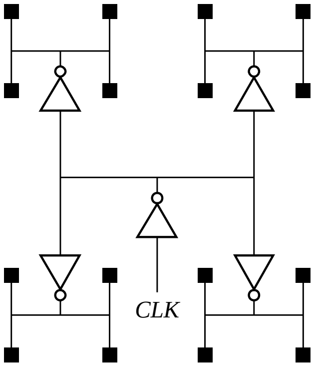

- Advantages
  - – Ideally zero-skew
  - – Can be low power (depending on skew requirements)
  - – Low area (silicon and wiring)
  - – CAD tool friendly (regular)
- Disadvantages
  - – Sensitive to process variations
  - – Local clocking loads are inherently non-uniform

Realistic H-tree

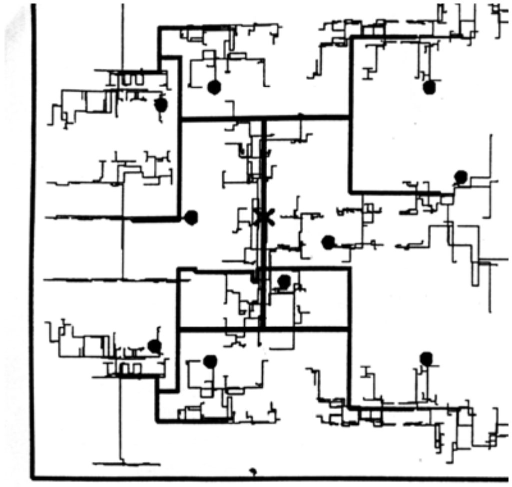

[Restle98]

Balancing a Tree

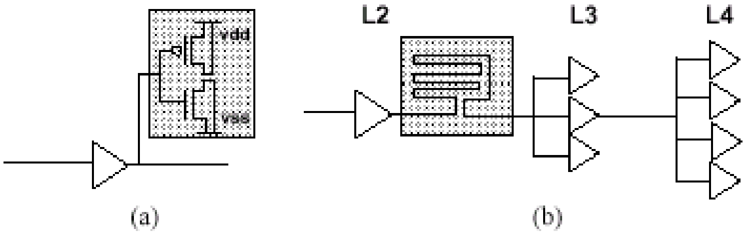

Con: Routing area often more valuable than silicon

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)

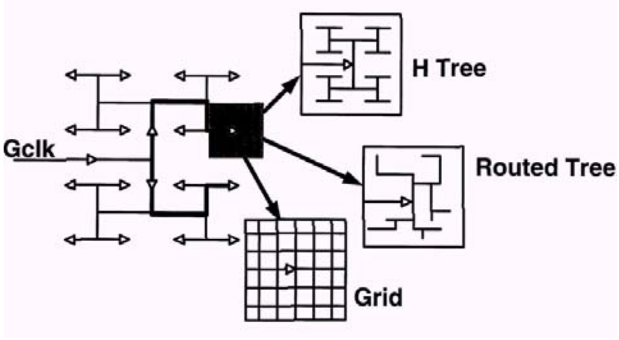

Summary

- 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

Power Distribution

- Power distribution network goals
  - Carry current from pads to transistors on chip
  - Maintain stable voltage with low noise
  - Provide average and peak power demands
  - Provide current return paths for signals
  - Avoid electromigration & self-heating wearout
  - Consume little chip and wire area
  - Easy to lay out

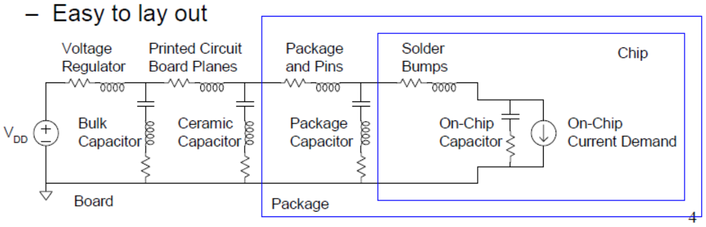

Stability Requirement

- $V_{DD} = V_{DDnominal} - V_{droop}$
- Want V_{droop} < ±10% of V_{DD}
- Sources of V_{droop}
  - IR drop
  - L di/dt noise
- I_{DD} changes on many time scales

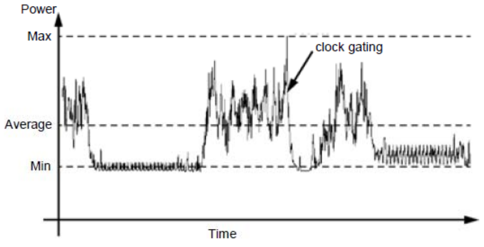

IR Drop

- Current supplied from voltage source to switching devices traverses a non-ideal power network. This causes the supply voltage appearing at the switching devices to deviate from ideal voltages

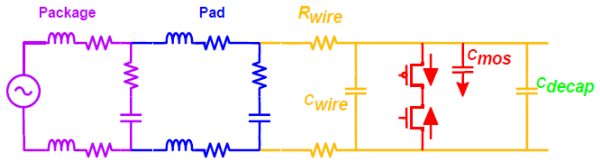

IR Drop Problem

- In response to inputs, gates switch drawing current from the power rails
- As current is drawn, the voltage supplied to gates deviates from ideal supply voltage

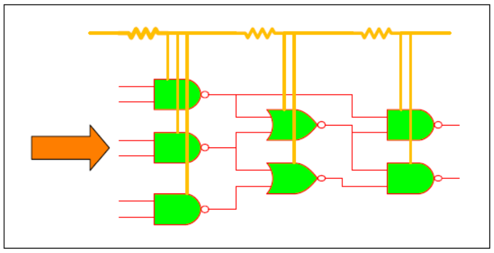

Technology Trends

- Power consumption of chips increasing while power supply voltages are being reduced
  - greater current demands and larger drops
- Chips have become larger
  - implies longer power lines and hence larger drops
- Impact of IR-drops on performance and signal integrity is stronger
  - IR-drops can become a larger percentage of rail-to-rail voltages

Analysis and correction of IR-drop integrity is critical

Power distribution networks will take up more of the metallization (and design) resources

Performance Degradation

- Reduction in rail-to-rail voltage degrades performance
  - Current drive
  - Propagation delay

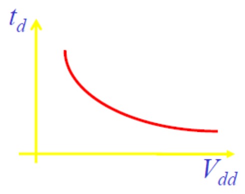

- Cell characterization and custom circuit design done with pre-defined budgets for IR-drop
  - drops larger than the budget causes performance problems

Reduced Noise Margins

- Noise margin reduces as rail-to-rail voltage swing reduces
  - Circuits become more susceptible to false switching due to noise

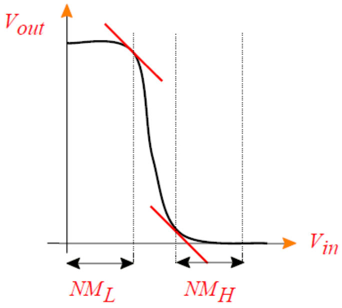

Dealing with IR Drop

- Large drop at a particular location
  - Widen existing wires
  - Add new wires
- Add more power pads

Electromigration

- Electromigration is the gradual wear of metal wires due to migration of metal atoms
  - – protrusions and holes can cause opens or shorts
  - – increased resistance causing higher voltage drops
- EM is directly proportional to DC current density
- High current density in Vdd and Gnd wires causes electromigration problems
  - – power lines are more vulnerable than signal lines because of unidirectional current

Electromigration

- Vias are especially susceptible to electromigration problems
  - – Typically a bottleneck in terms of cross-sectional area through which current can flow
  - – Use via arrays
  - – Current crowding stresses vias unequally
    - Current crowds around edges of vias rather than using the entire cross-sectional area

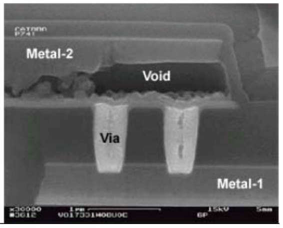

Ldi/dt

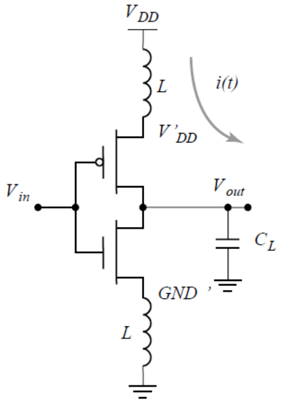

Impact of inductance on supply voltages:
- Change in current induces a change in voltage
- Longer supply lines and cheaper packages have larger *L*

Dealing with Ldi/dt

- Separate power pins for I/O pads and chip core
- Multiple power and ground pins
- Increase rise and fall times of off-chip signals to maximum extent allowable
- Use advanced packaging technologies
- Add decoupling capacitances
  - Typically use MOSFETs with source/drain shorted, large capacitance per unit area due to thin gate oxide

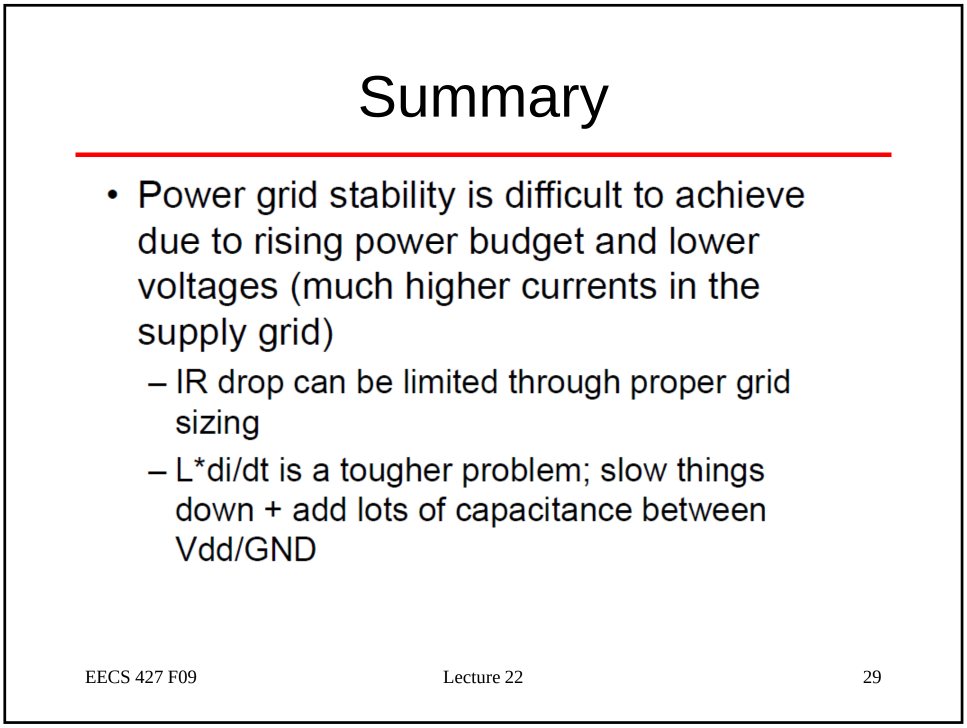

Summary

- Power grid stability is difficult to achieve due to rising power budget and lower voltages (much higher currents in the supply grid)
  - IR drop can be limited through proper grid sizing
  - L*di/dt is a tougher problem; slow things down + add lots of capacitance between Vdd/GND