VLSI Power Delivery

David Ayers

Outline

- Die power delivery
  - Die power goals
  - Typical processor power grid
  - Transistor power noise
  - SSN noise and control (decoupling)
  - Package connections
  - Large di/dt's
- System power delivery
  - Compents: VR's; MB; Socket; Package
  - Capacitor arrays
  - Freqeuncy analysis and resonance
- Related topics
  - Signal return path and cross-talk
  - IO power delivery
  - Filtered supplies for sensitive circuits
  - Scaling

Die Power Delivery

Goals of the Die Power Network

- Do's
  - Deliver power from the package to the transistors with little voltage drop
  - Deliver charge from die capacitances to transistors to control noise spikes
  - Provide signal return and shielding
- Don'ts
  - Network should not wear out from electromigration and self-heating
  - No onerous layout requirements
  - Area usage should be minimized
- Designer must balance the competing objectives
  - For example, small voltage drop competes with minimizing area (metal) usage
- Typical solution has a regular grid in the upper layers

6 Layer Power Grid Example -- CBD

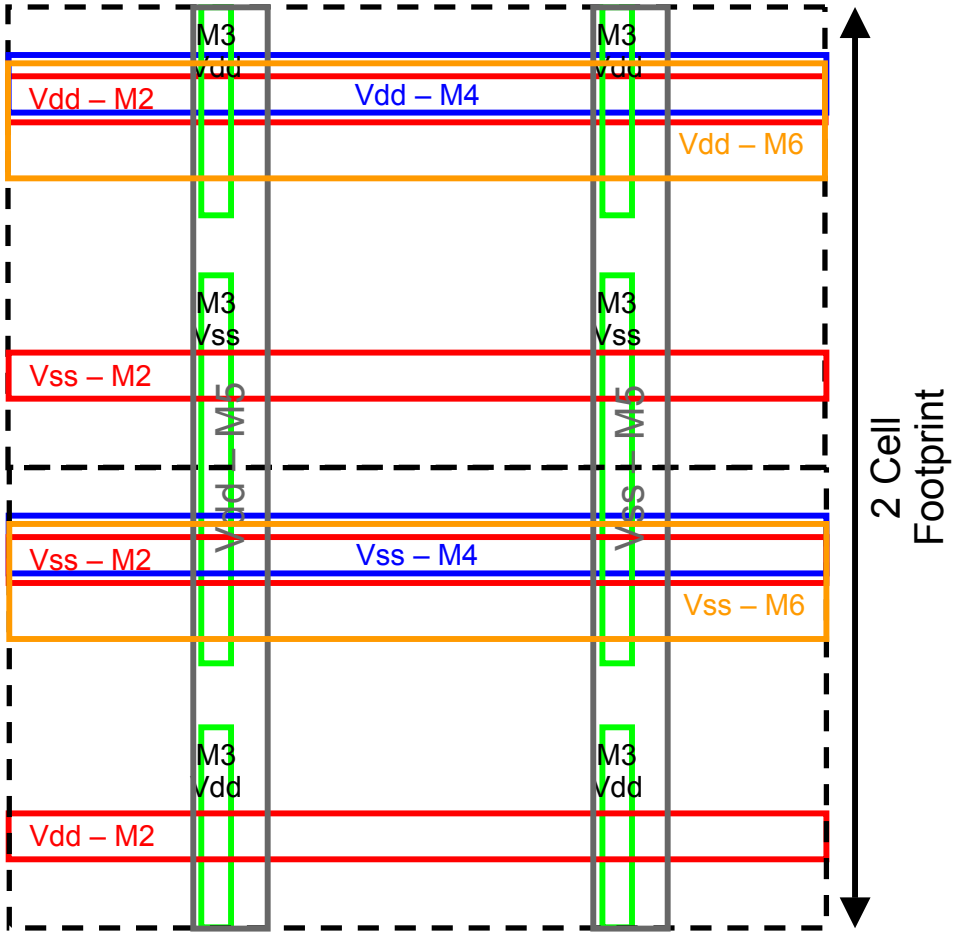

- Representative power grid design for 6 layer CBD shown
  - Custom layout may not be as regular at M2 & M3
- M2 is mirrored for well abutment
- M3 power shares tracks to limit metal usage and increase via counts
- Vias located at all next layer crossings
- Power metals are stacked as much as practical to simplify via stacks
  - Provides short access to thick upper layers
  - Thick and wide upper layers dominate the structure

Local Layout Considerations

- Transistor technology trends
  - Trend towards single poly direction for Optical Proximity Correction and Phase Shift Mask generation
  - Leg length being limited by increased poly resistance
    - Though fully silicided gates are being reported [3], [4]
  - Simplifies power layout and limits M1 leg lengths
    - Makes full gridding more practical
  - Vertical flow becomes very important, even dominate
    - Must consider vias in chip power models
  - High current density driving the need for high via counts for EM and voltage drop
    - May need more or wider power stripes to accommodate vias
    - Via counts should be in proportion to via resistance
      - Example: V5 = 1 Ohm, V1 = 3 Ohms ➔ should have 3x as many V1's
    - May need to push for larger size power vias

Transistor Behavior

- Graph of current vs. time from simulation of a repeater segment in 90 nm technology
- Behavior and local decoupling need is independent of clock frequency!
- Charge needed can be computed by integrating the current waveform
- Decoupling capacitance needed can be estimated by Q/ΔV
  - ΔV is the budgeted voltage drop

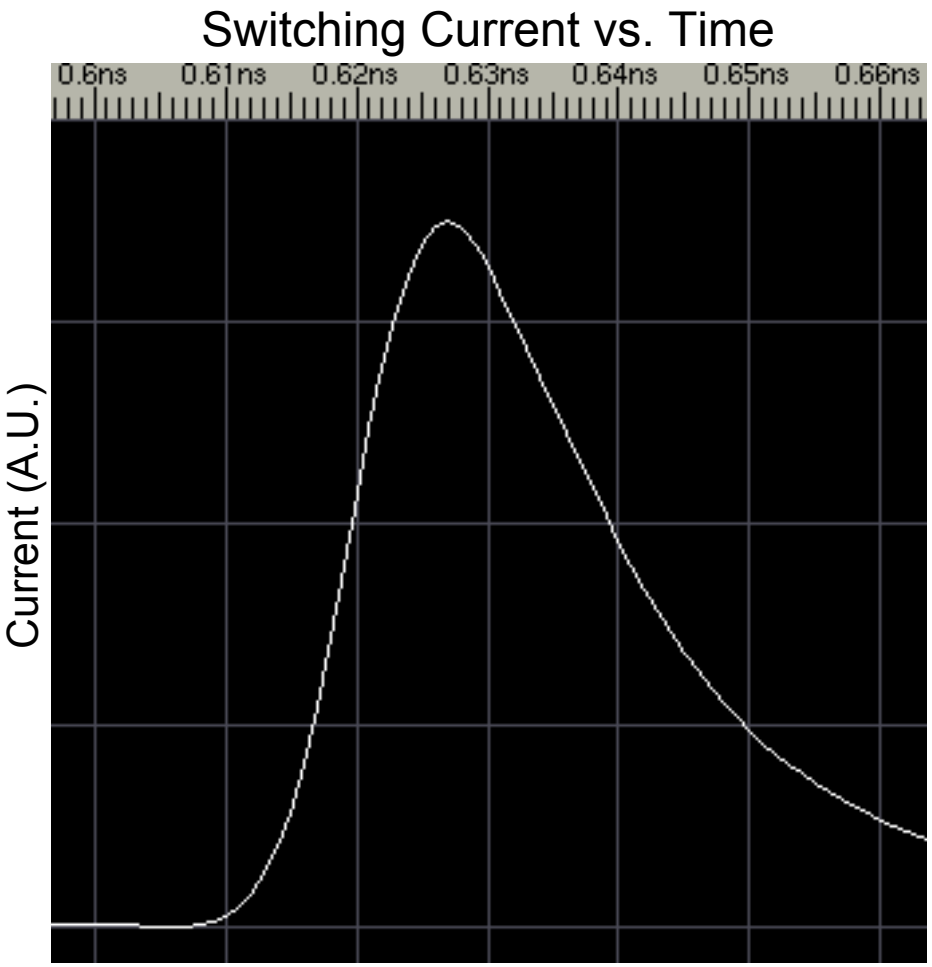

White (middle wave) = moderate load, FO ~ 4

Transistor Behavior (Cont'd)

- The worst IR drop point is at the peak
  - Worst Ldi/dt drop is at the inflection point halfway up the ramp
- The rising edge is nearly identical regardless of the load!
- Note that the transistor current is very fast -- peaks in < 20 ps
- Decoupling must be located close enough to be reached in time
  - Must beat current peak
- Package capacitors way too slow (nanosecond tau's)

Switching Current vs. Time

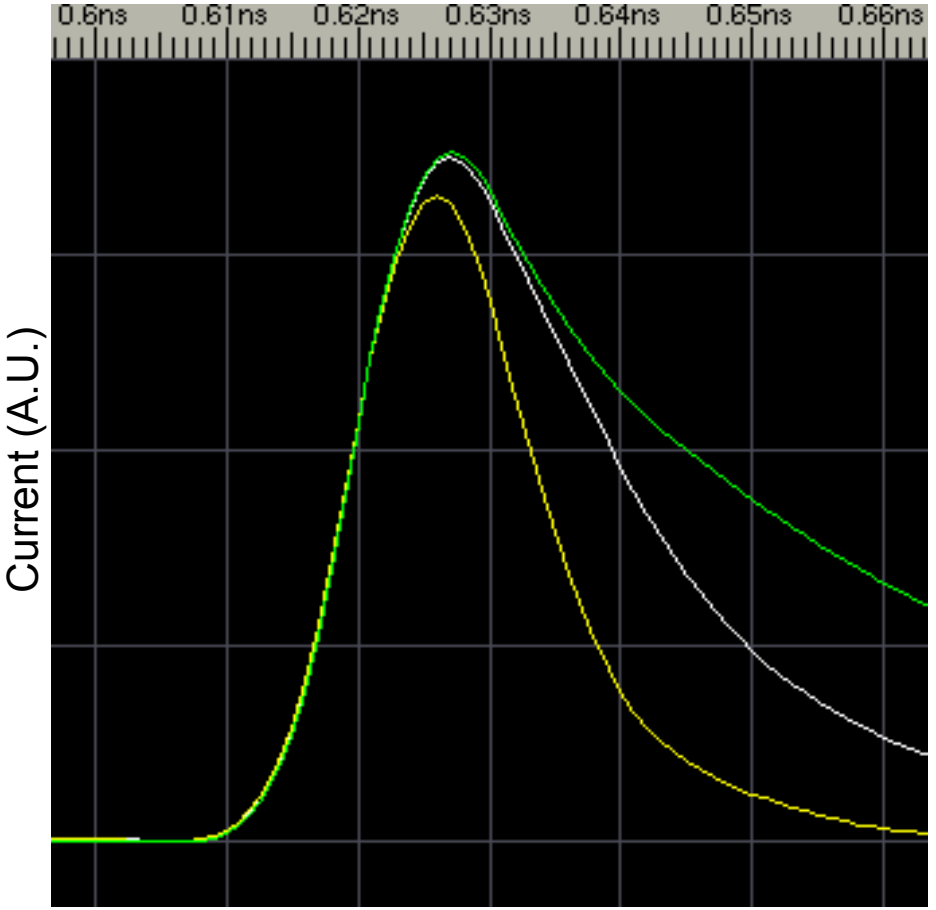

White (middle wave) = moderate load, FO ~ 4
Green (upper wave) = heavy load, FO ~ 7
Yellow (lower wave) = light load, FO ~ 2.5

Impact on Nearby Logic

- Very fast transistor switching means very fast noise spikes
- Random block of logic is usually not a big noise concern
  - Thousands of scattered small transistors fire at various times in a clock cycle
    - Not enough 'stuff' firing at once to cause a serious disturbance
    - More like white noise
  - Bad case will be bank of synchronous drivers (like repeaters)
    - 64-256 large drivers firing synchronously
- Wave shown is from a power model repeater bank simulation with 90 nm technology
  - Spike droops up to 19% of Vdd
  - But droop only exceeds 5% of Vdd for < 25 ps

Net Voltage (Vdd-Vss) vs. Time

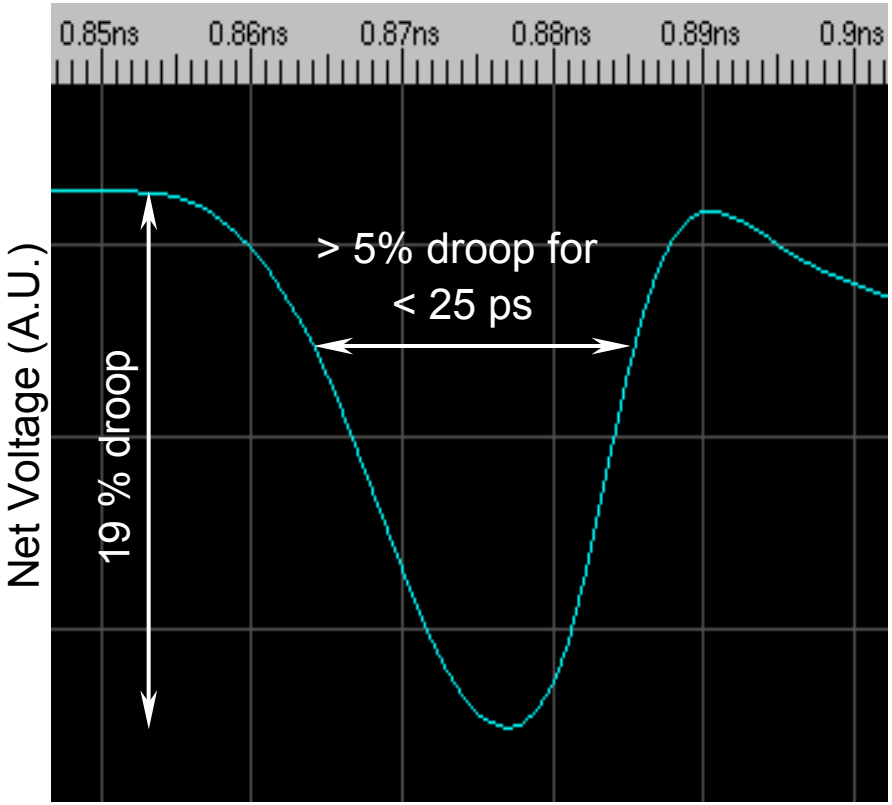

- With a clock cycle > 200 ps, there is minimal delay impact to nearby logic from one spike
  - Is extra decoupling really needed?
  - Noise spikes have the greatest speed impact on the repeated signal itself

Droop vs. Decap Distance and Die Metal

- Simulations from 180 nm technology node
  - Capacitors placed at various distances from noise source
- Note noise increase as capacitors are placed further away
- Substantial improvement with increasing power metal use

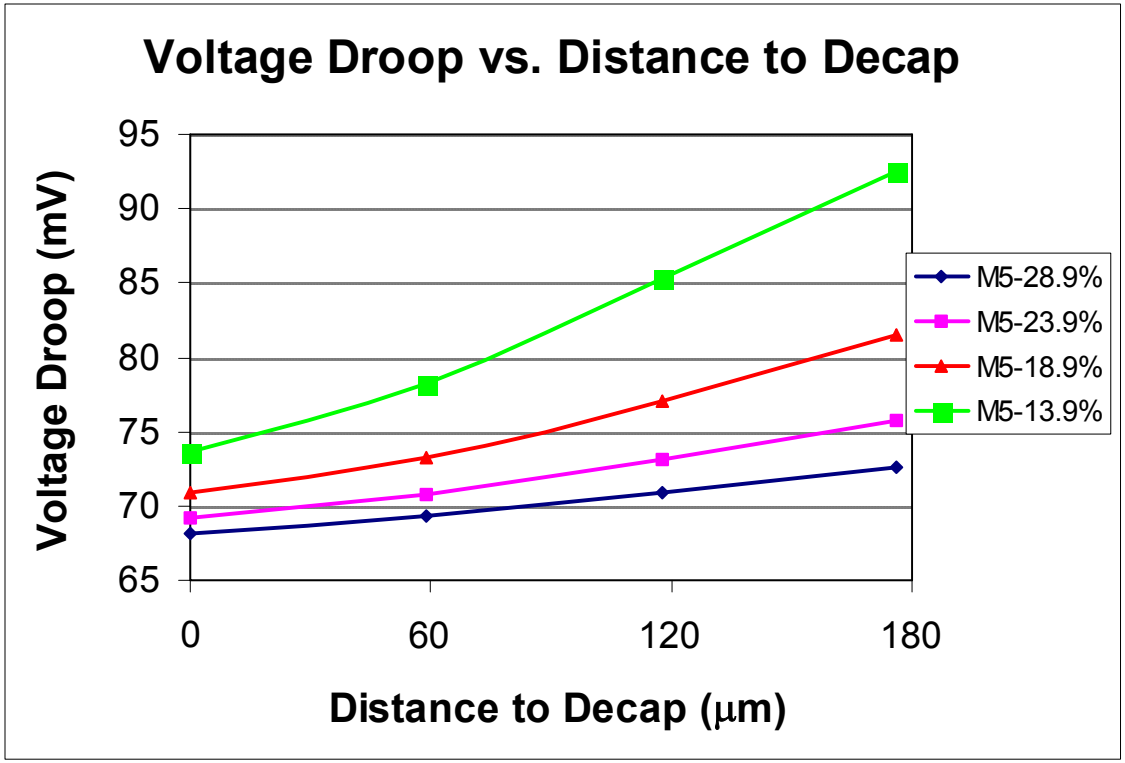

Grid Propagation

- Die propagation is slow since resistance dominates
  - Typical top layer resistivity of 25-30 mOhms/sq
    - About 6-8x that (~ 200 mOhms/sq) for a single direction if half the layer is used for power wires
      - Factor of 2 each for half of a layer, half to Vdd or Vss, half line and space
  - Typical circuit capacitance density is ~ 1nF/mm^2 (90 nm node)
  - RC time constant is ~ 200 ps/mm one way!
  - Most droops are in the 10-20% range (not 1/e = 63% droop for an RC time constant)
    - Thus 10-20% droop propagates in the ~30 ps/mm range
    - Need to travel both to and from capacitance
    - Plus there is non-quasi static delay in decoupling cells of several ps
- This limits the useful distance at which decoupling capacitors can be placed to a few hundred um or less
    - 30 ps/mm limits decoupling distance to ~200 um to respond to a 20 ps current spike

Capacitance Density

- Package planes have 30-100 um separation
  - < 0.001 fF/um^2
- Die metals have ~ 0.5 um of separation
  - ~ 0.07 fF/um^2 (about 4x this value for 8 metal layers)
- MOS cap has ~ 2 nm of separation (90 nm node)
  - ~ 18 fF/um^2
- Current MIM (metal-insulator-metal) capacitor technologies are reaching ~ 1 fF/um^2 (see the summary table in [1])
  - 1 fF/um^2 is probably not very useful for bulk decoupling, need about 10x that
  - A higher density with a single mask was recently reported [2]
- MOS type caps remain as the main source of supplemental decap on die at 90 nm
  - Leakage is limiting usefulness – may need special structures

Decoupling Capacitor Design

- Want high capacitance density and low resistance for fast response
  - High density achieved by maximizing the poly gate oxide area
  - Low resistance achieved by limiting the distance between contacts to ~ 1um
- Decoupling added for global di/dt changes (1st droop) can have longer distance between contacts

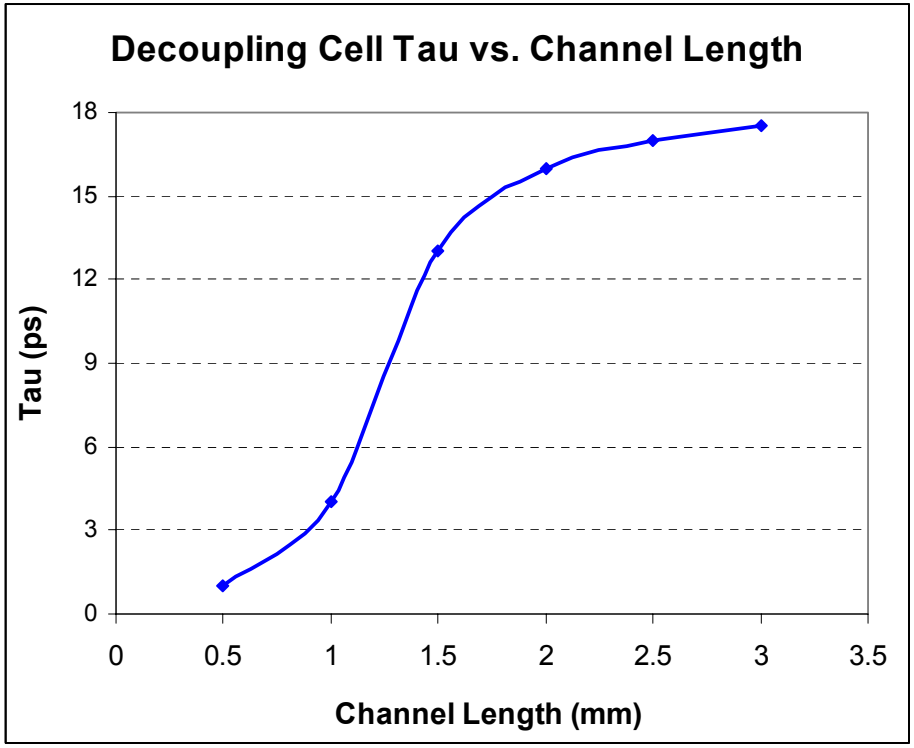

NMOS Transistor Style Capacitor Simulated in 130 nm Technology

Decoupling Capacitor Design (Cont'd)

- Cell type can be important
  - NMOS faster than PMOS inversion cells
  - PMOS accumulation cells can be faster than inversion but require wells which eat up space
  - Gate oxide leakage concerns may force accumulation cells
    - Work function shift reduces leakage
    - But capacitance rolls off at lower voltages (see graph)
    - Not well suited for analog circuit applications

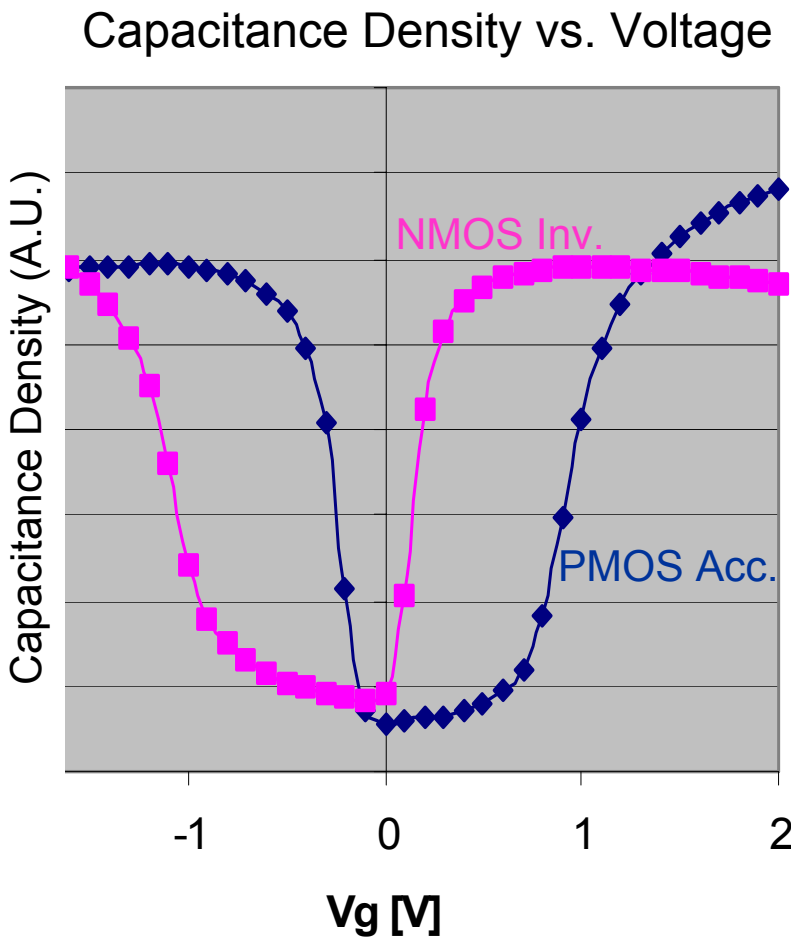

Dense Capacitor Layout Example

- NMOS ‘waffle’ type layout shown
  - Poly in green, M1 in red
- Essentially a sheet of poly
  - Non-minimum openings for the silicon contacts (Vss)
  - Field oxide bumps for the poly contacts (Vcc)
- Achieves very high capacitance density with good contact spacing for low resistance and tau

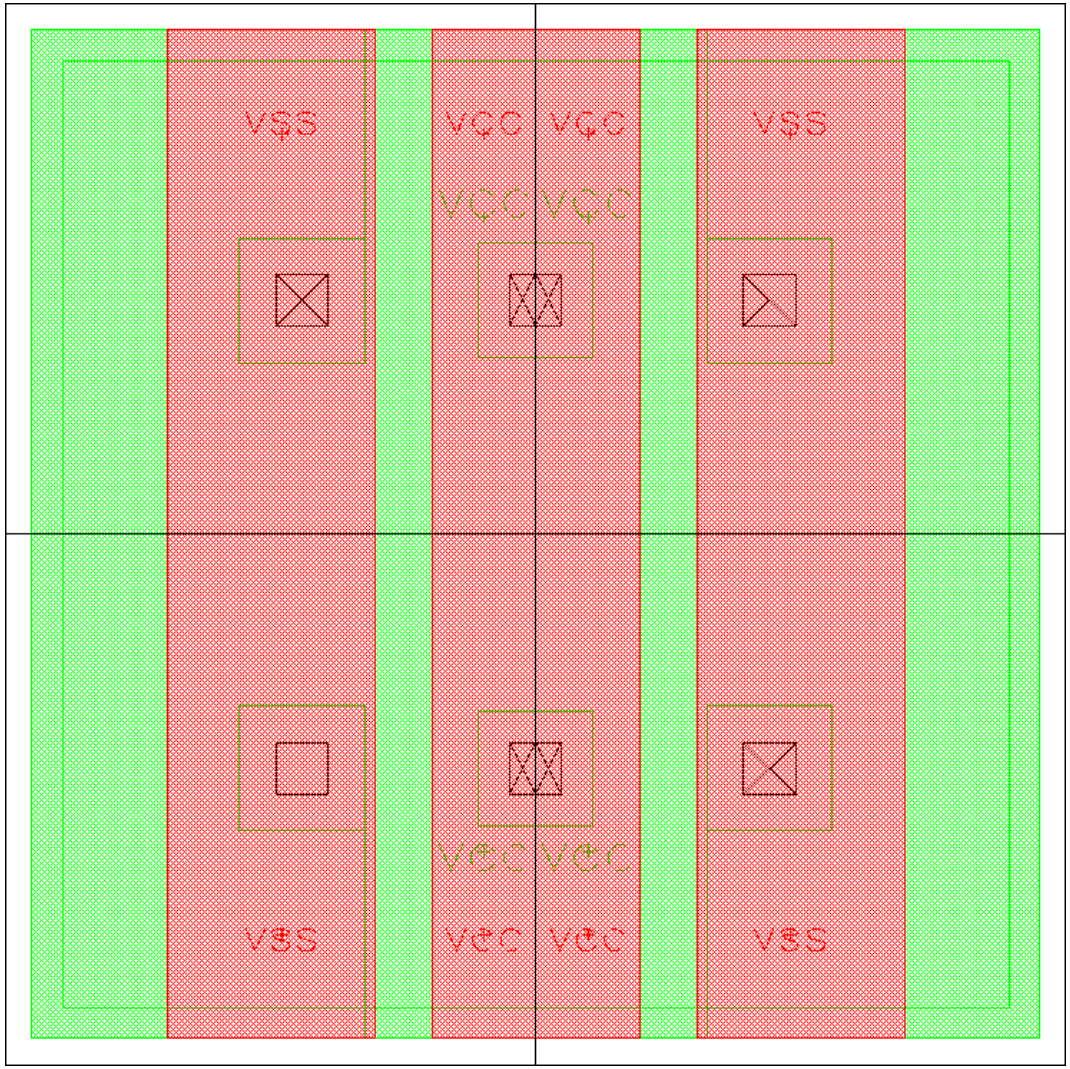

Decoupling Capacitor Design (Cont'd)

- Must be aware of defects and planarity impacts
  - High poly and M1 density may increase the variation in nearby devices
    - Lower poly density means less capacitance per unit of area
    - Need to make trade-offs
  - May have millions of cells
    - Use greater than minimum spacing to reduce defect risk
- Need unit cells which can be easily built into arrays by tools

Package Connection

- C4 bump pitch has not been scaling as fast as transistor technology while current density is scaling
  - Result is increasing current per bump which will stretch reliability limits
- Note that only a few small areas have the highest current
  - Technology and uarch solutions are likely to be needed
- Increased top and second layer metal resources will also be needed

C4 Bump Current Density for a Processor

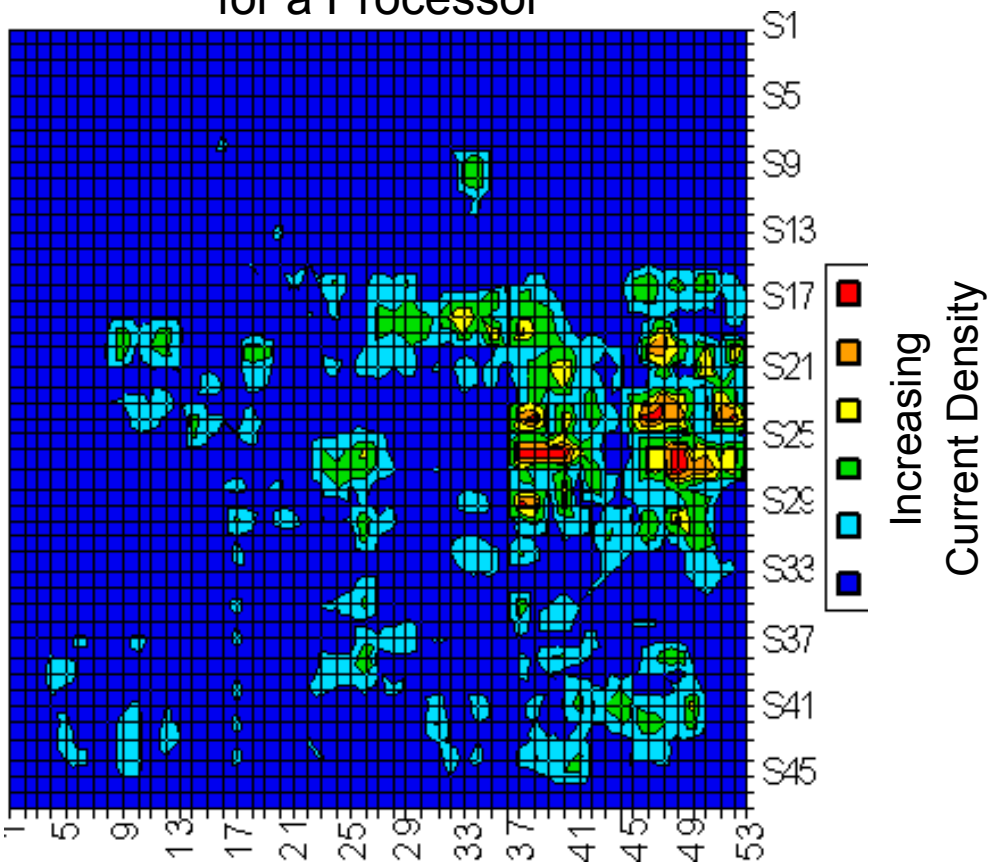

Large di/dt Swings

- So far we have mostly discussed local die noise
- Today's processors use extensive clock gating to reduce power consumption
  - Clocks and clocked elements consume about 50% of active (non-leakage) power
  - A largely inactive processor can have very low active power consumption
    - Can be less than 50% of peak power
- Processors can transition from a low power state to a peak power state as fast as the pipelines can fill
  - For the 90 nm generation, this can be less than 20 cycles (<5 ns) for some processors
  - Processors are approaching 100 W
  - Can have di's as high as 50 A
- Since leakage prevents much decap from being added, such swings will overwhelm die decap very quickly

System Power Delivery

Typical Power Delivery System

- 2 processor MB design shown
- VR current brought in to processors on ~2 sides
- Note the levels of decoupling
  1. Die (MOS)
  2. Back of package
  3. High speed MB
  4. Low speed MB
- VR current brought in to processors on 2 sides to reduce impedance
- VR located close to processor

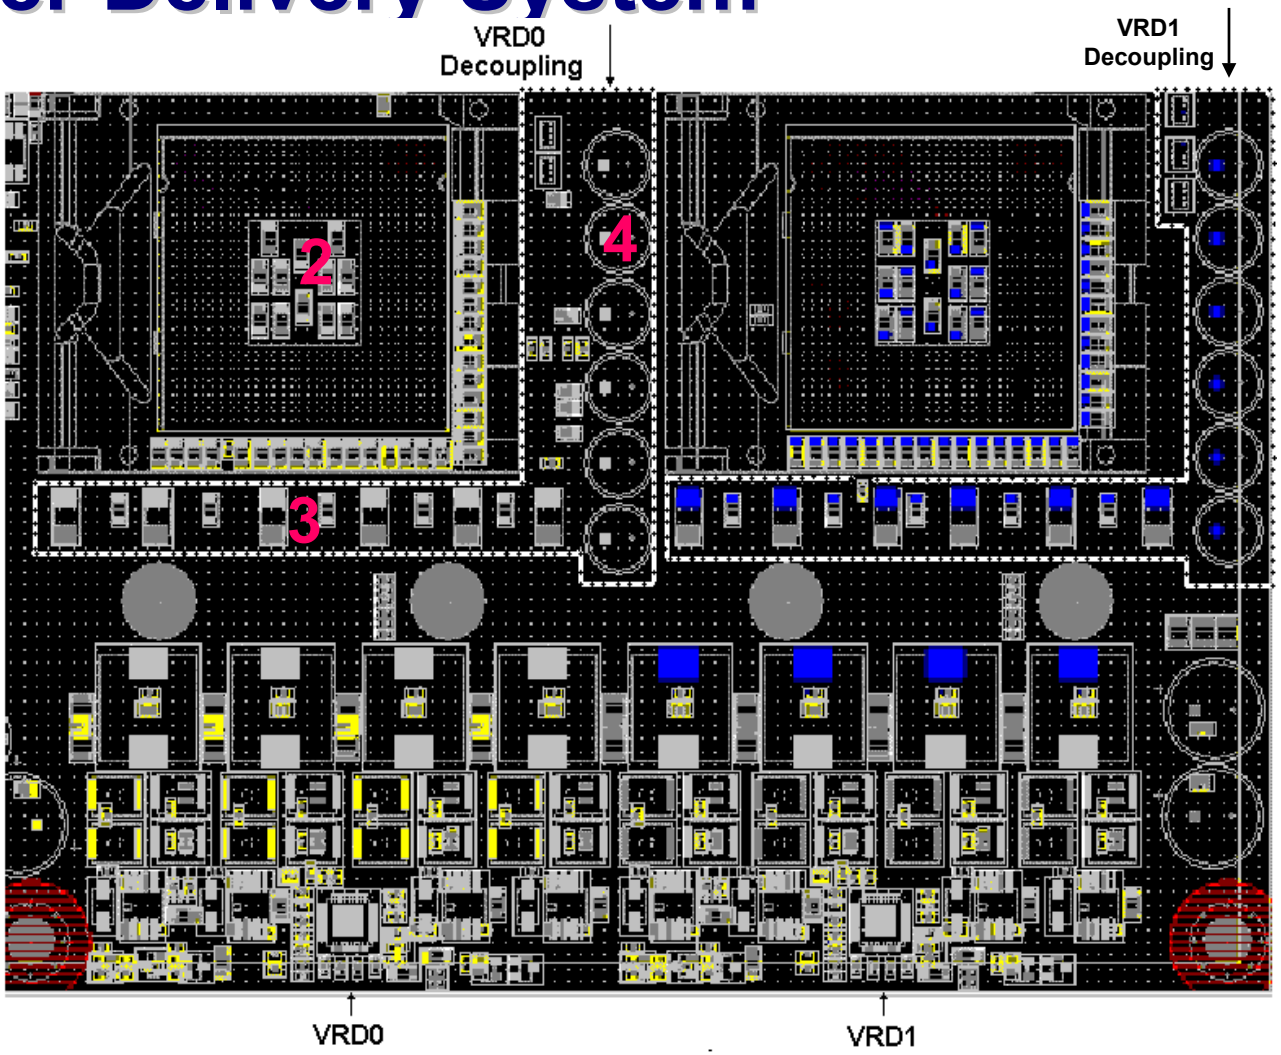

Packaging Cross-Section

- A sample processor cross-section is shown below
  - May or may not have a heat spreader
  - May have die side capacitors as well as land side
  - Package may have 4-14 layers depending on number of signals and cost structure of market (low-end desktop to high-end server)
  - May have an additional layer of package (interposer) for space transformation and for housing additional components
- Power must penetrate through the socket and package

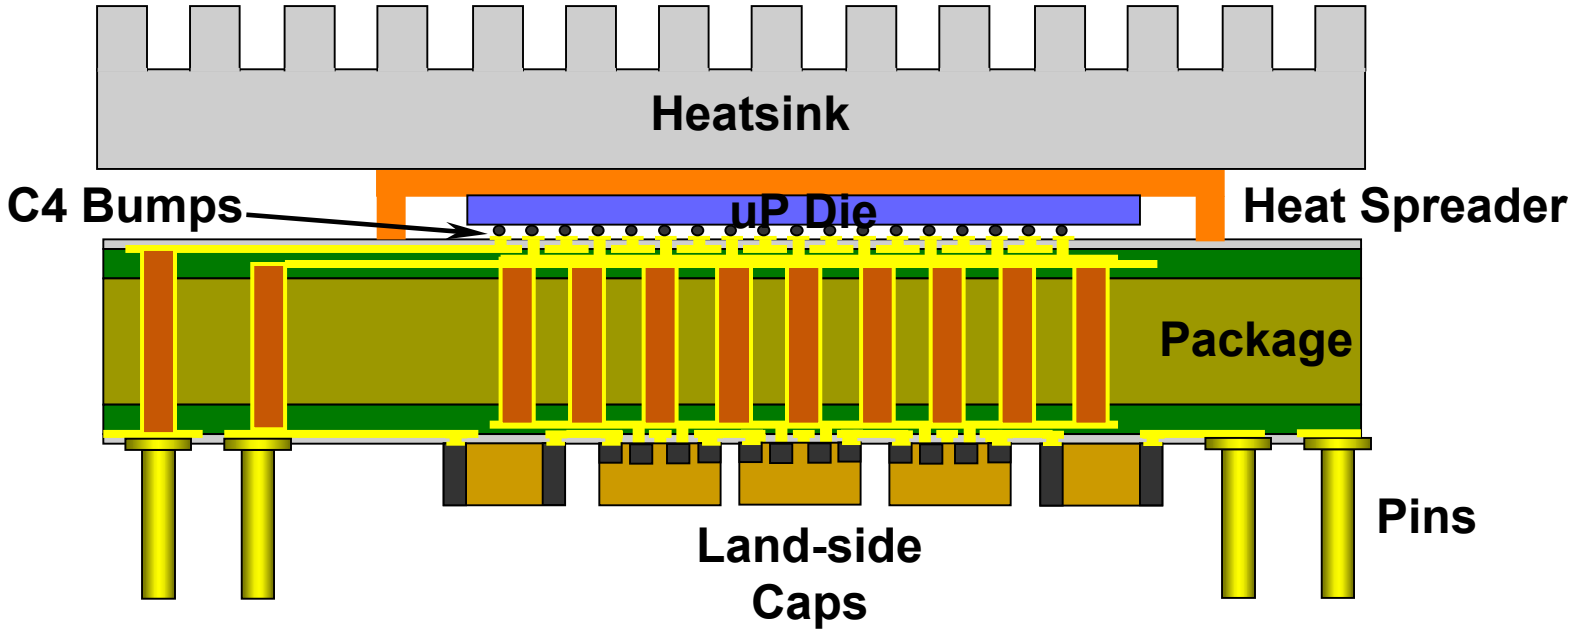

Factors in Determining Decoupling

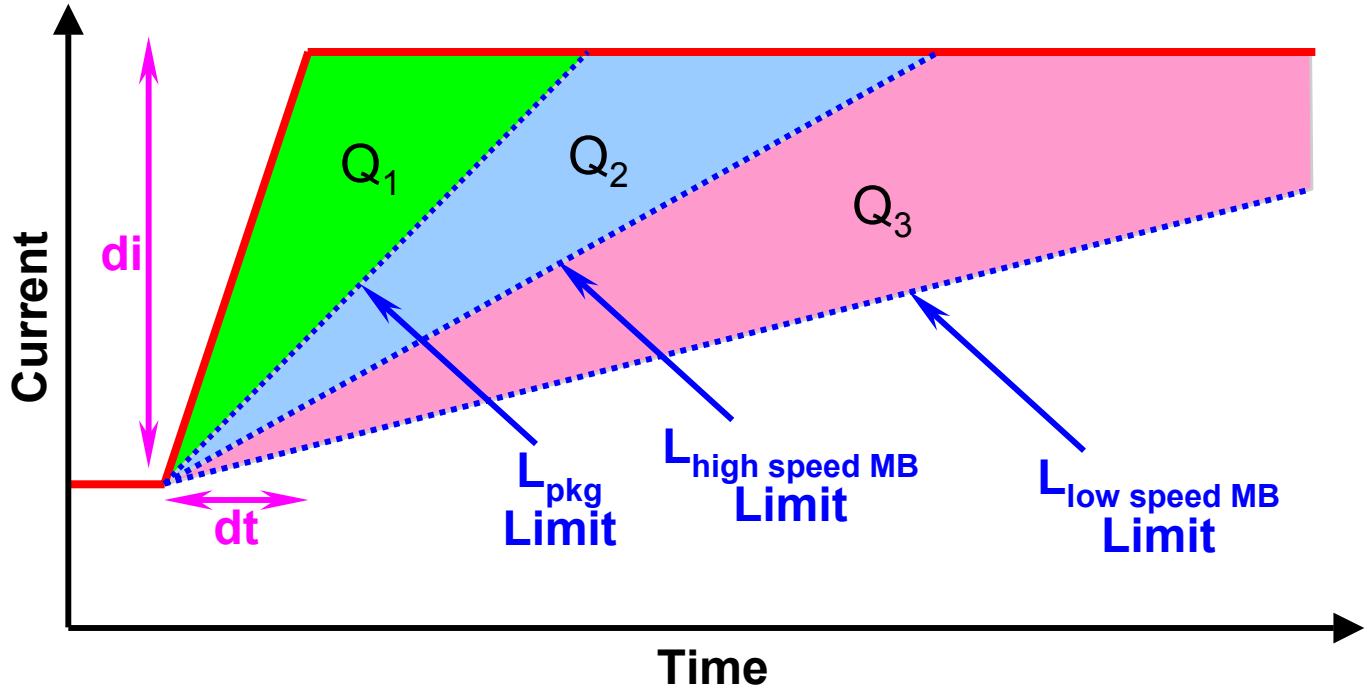

- The area of triangle Q_1 determines the need for die capacitance
  - $C_{die} = Q_1 / \Delta V$; determined by di, dt, L_{pkg}, and the voltage drop target
- The area of triangle Q_2 determines the need for package capacitance
  - $C_{pkg} = Q_2 / \Delta V$; determined by di, L_{pkg}, L_{HSMB}, and the voltage drop target
- The area of triangle Q_3 determines the need for board capacitance
  - $C_{board} = Q_3 / \Delta V$; determined by di, L_{HSMB}, L_{LSMB}, and the voltage drop target

Power Delivery Implications -- dt

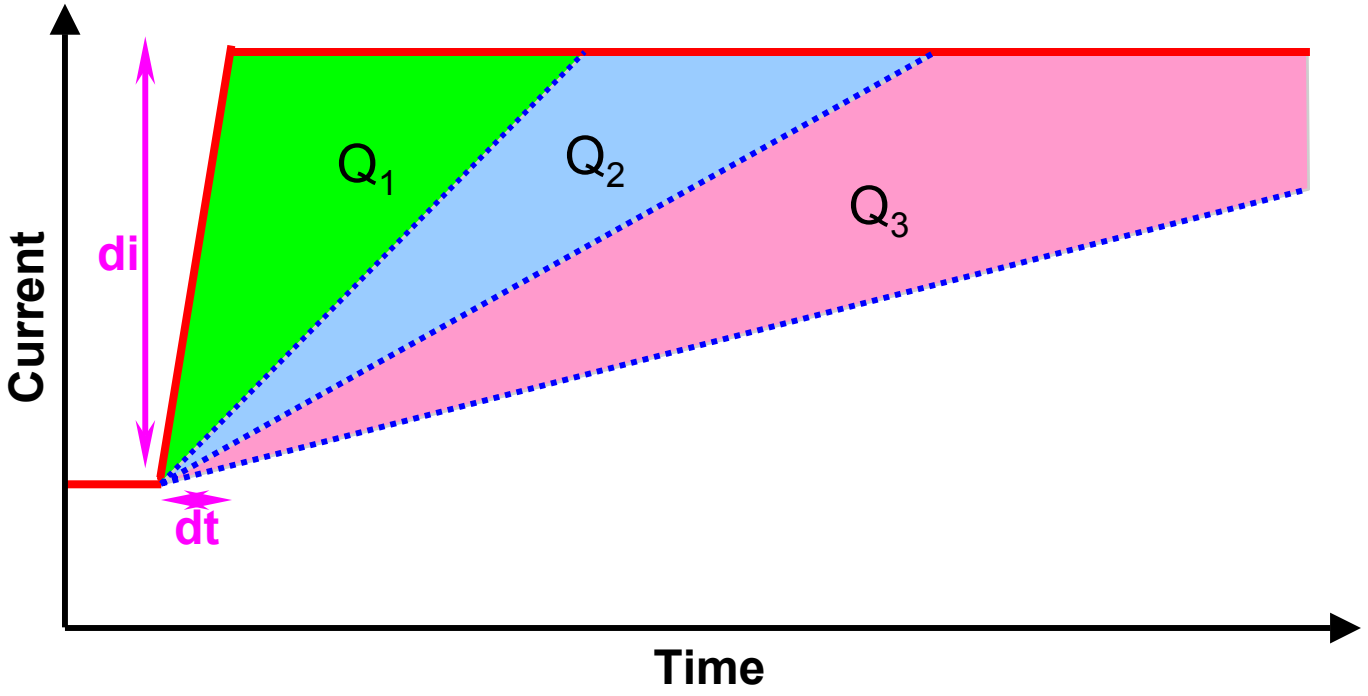

- Picture shows dt decreased by 2x from previous page -- small impact
- Capacitances are proportional to triangle areas
  - Note that the area of the Q_1 triangle (die capacitance) increases by less than 2x
  - Area of the other triangles (other capacitors) are unaffected
  - See [6] for a treatment of di/dt control

Power Delivery Implications -- Imax

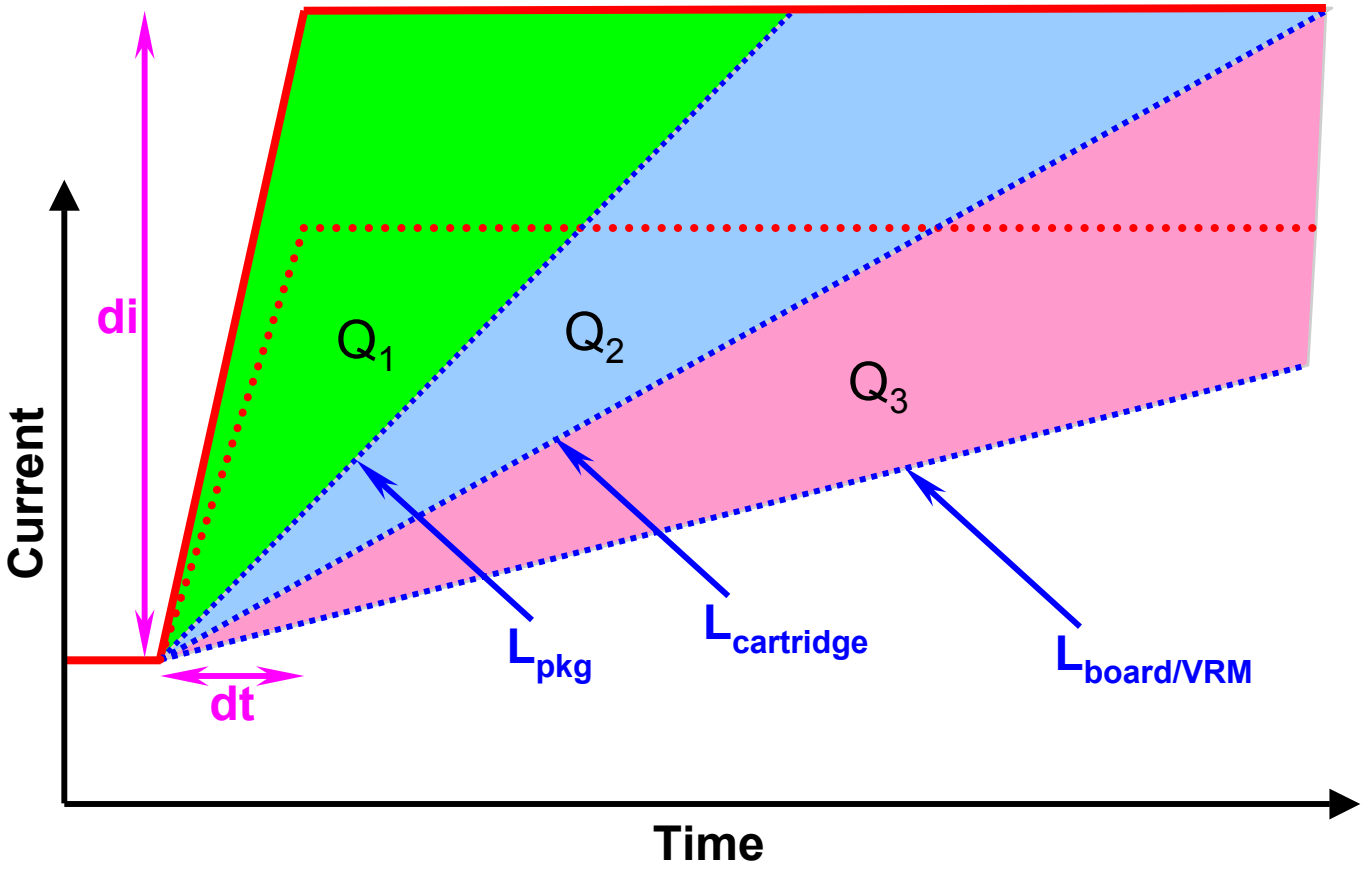

- An increase in di has a big impact on all the capacitances each of which is proportional to the triangle areas
  - Square relation for area: 2x increase in di increases the triangles by 4x!
  - Even greater increase for Q_1
- Reducing di is most effective for voltage control

Capacitor Current Example – 4 Levels

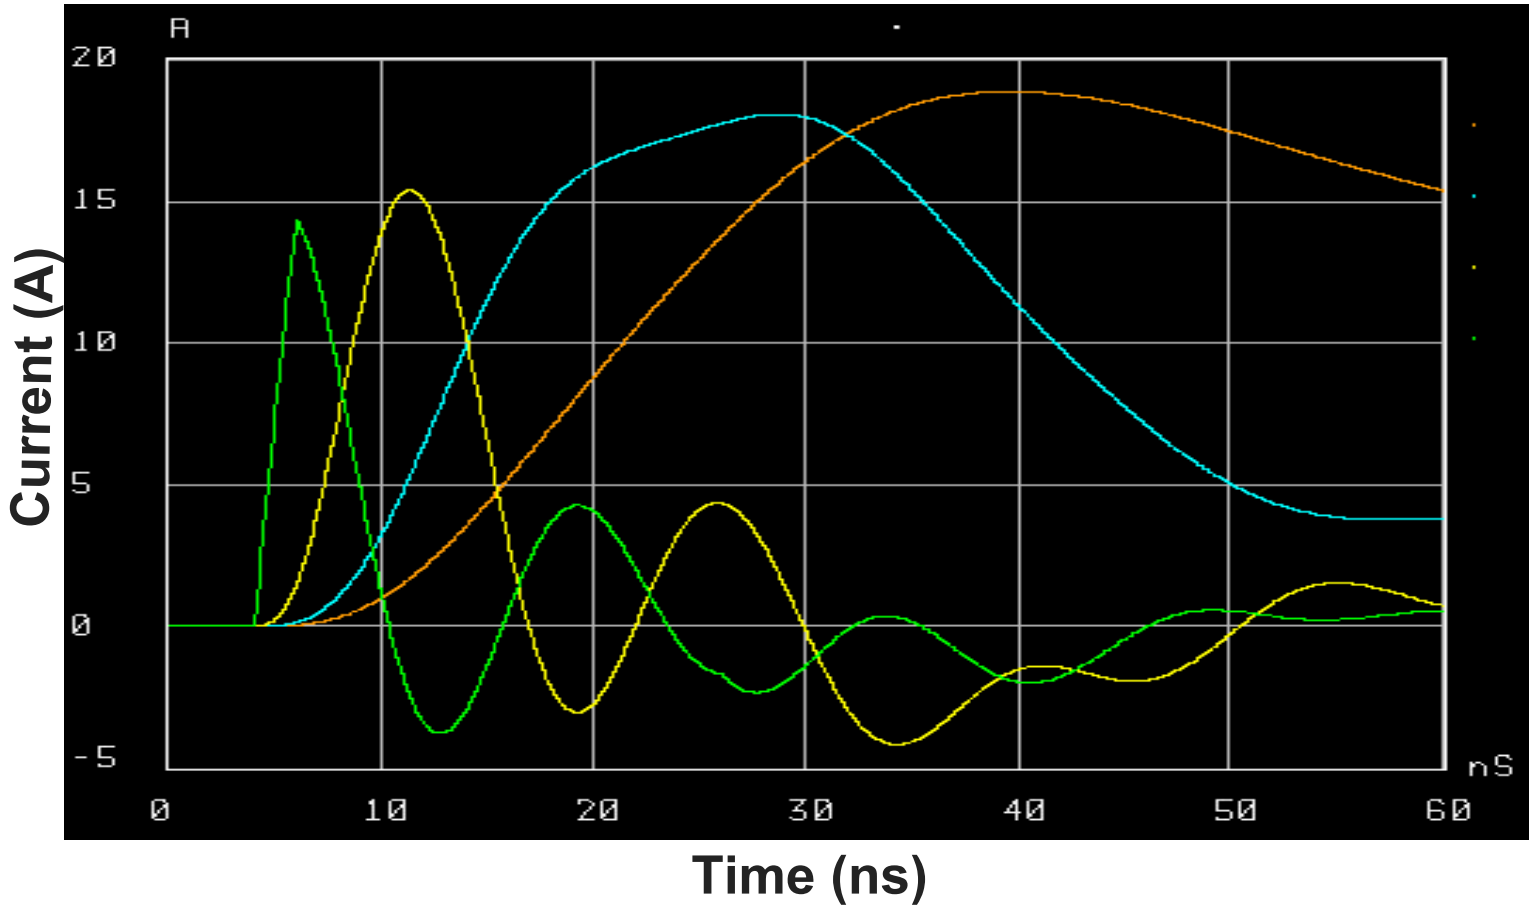

- Local voltage minima will be reached at each cross-over point
- Ideal design would have these minima equal
  - More cap levels would allow us to approach a flat impedance vs. frequency curve

Time Response

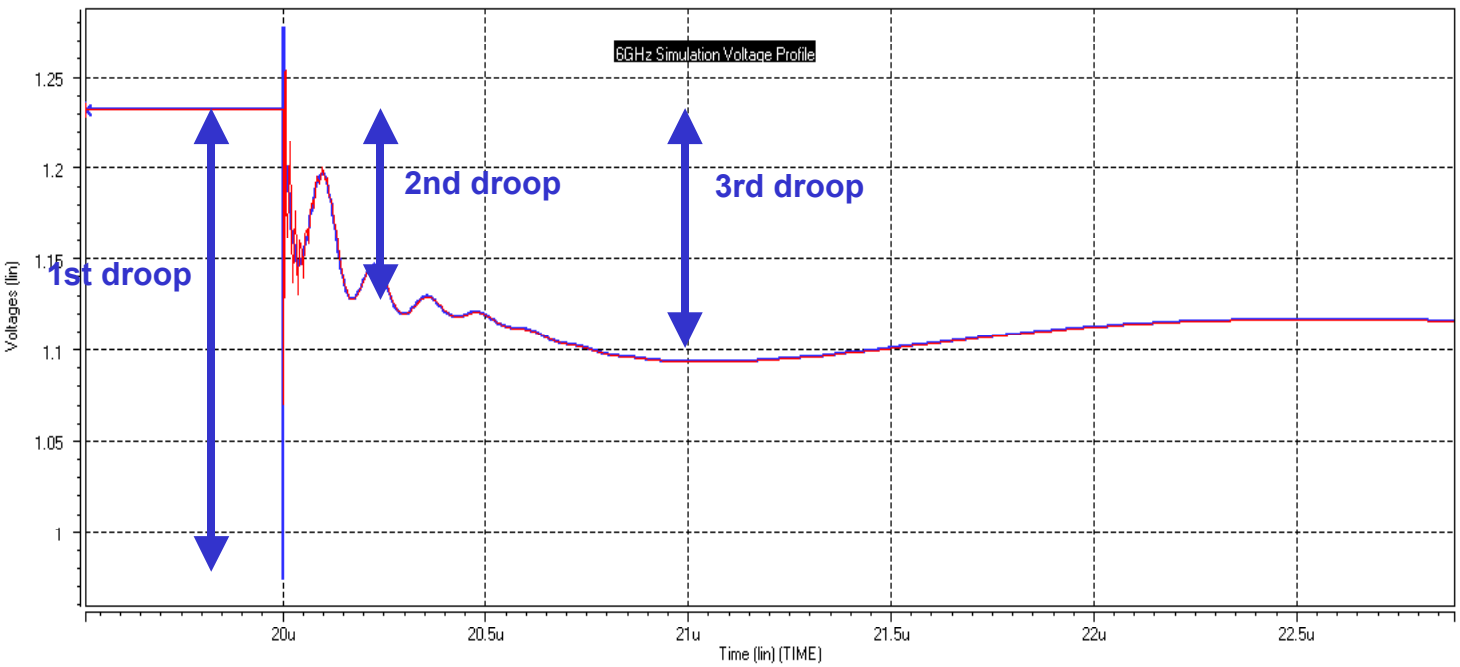

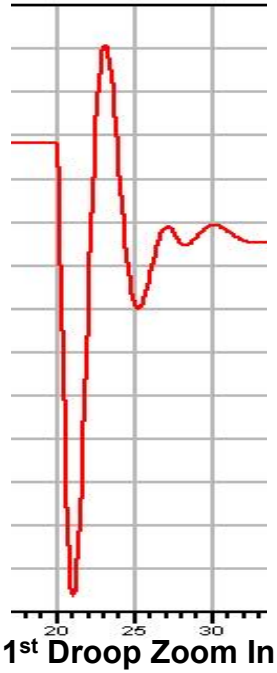

1st Droop Zoom In

- Graphs shows a simulated voltage response for a system with die, package, board, and VR (bulk) capacitor arrays
- One droop for each cross-over in source of current
  - Example: 1st droop reached when main source passes from die cap to package cap

Frequency Domain System Modeling

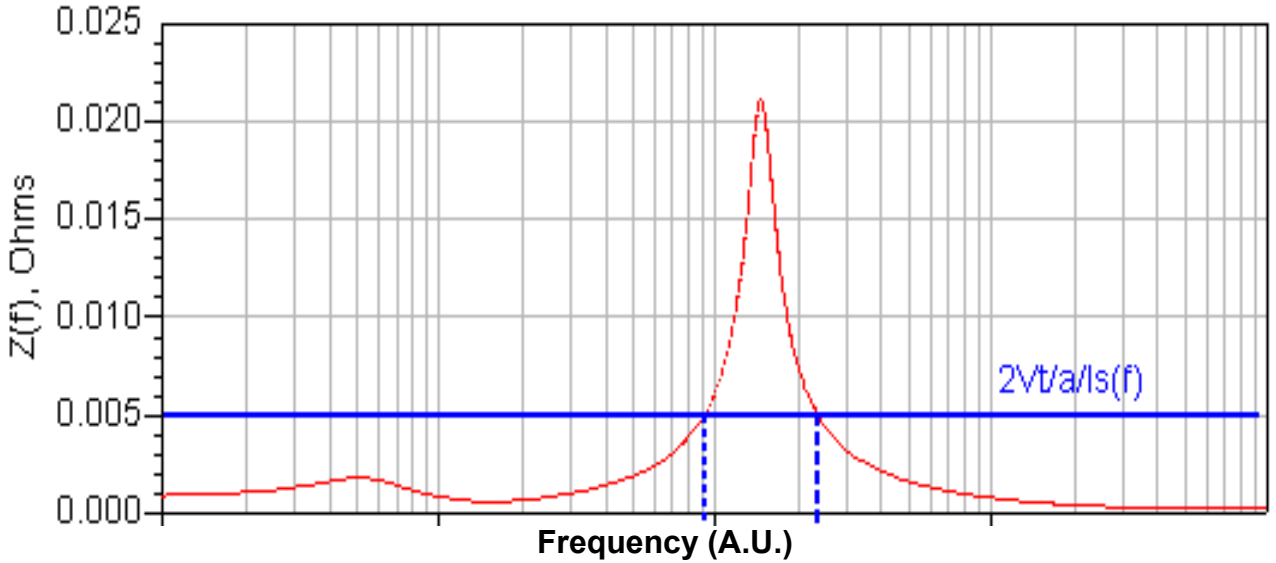

- Graph of power delivery impedance vs. frequency shown
  - Peak at the resonant frequency of the package/die system
  - Up slope caused by an inductance (ωL)
  - Down slope results from a capacitor ($1/\omega C$)
- Want to move the up slope right (by reducing the inductance)
- Or, move the down slope left (by increasing the capacitance
  - May be at odds with cost targets!
- Otherwise, may need special techniques to overcome or live with the resonance
  - Clamps, charged pumped capacitors, etc., [5]

Related Topics

Global Layout and Signal Return

- Power grid also supplies return path for signals
- Wide busses (~128b) can switch over an Amp of peak current
- Without adequate return path, inductive noise spikes can disturb signals
  - Additive to capacitive noise for certain patterns
- May need additional power wires to control inductive cross-talk
  - Example: half-shielded scheme shown below
- Robust thick upper metal power gridding can provide return path for lower layer signals
  - Can focus on signal density on lower layers

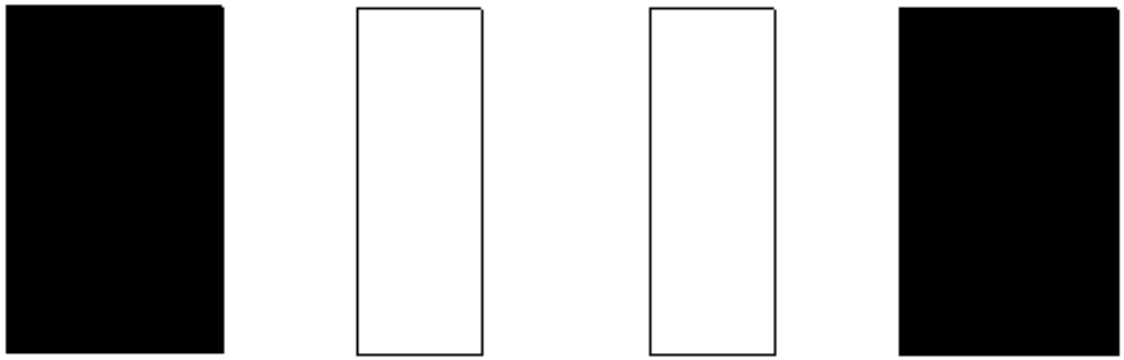

IO Power Delivery

- IO supplies (Vtt) are often separated from core power
  - Buffers on processor and chipset run off the IO supply
  - Signals usually referenced to Vss as much as practical
    - Avoids needing to decouple *between* supplies to establish a low impedance return path
- IO supply will need adequate decap and routing
  - Similar analysis to core power
  - Will also need frequency analysis
  - SSN analysis especially important
  - Will need partial (or full) package planes for Vtt
  - May need full Vss planes to shield the IO signals

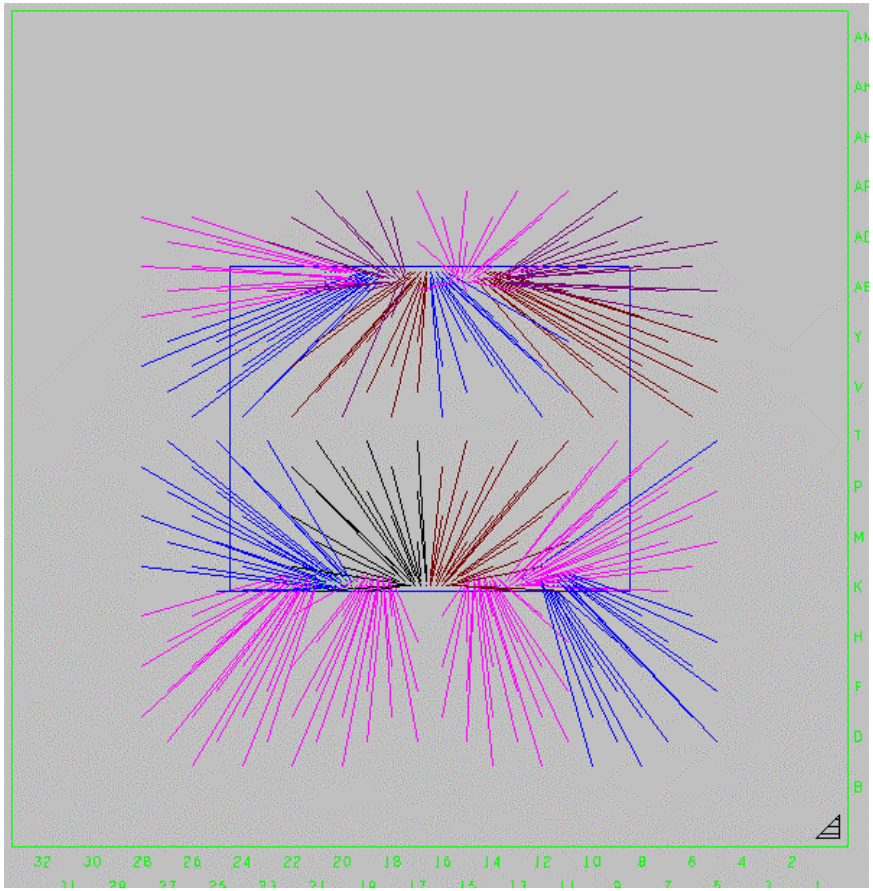

Example of Package IO Routing (Crow's Flights)

IO Power Issues

- Simplified models (like the one shown) have a path to ideal ground
  - This is non-physical
- Models must reflect true path including signal return path in the Vss network
  - Only way to properly reflect the interaction of Vdd (core supply) and Vtt (IO supply)
  - IO signaling will inject noise into the core Vss (and vice-versa)
- SSN will probably need control
  - Edge rate control
  - Die and package Vtt decoupling
  - Careful via and bump layout to keep inductive loops small
  - Data inversion to avoid more than 50% 0-1 or 1-0 transitions

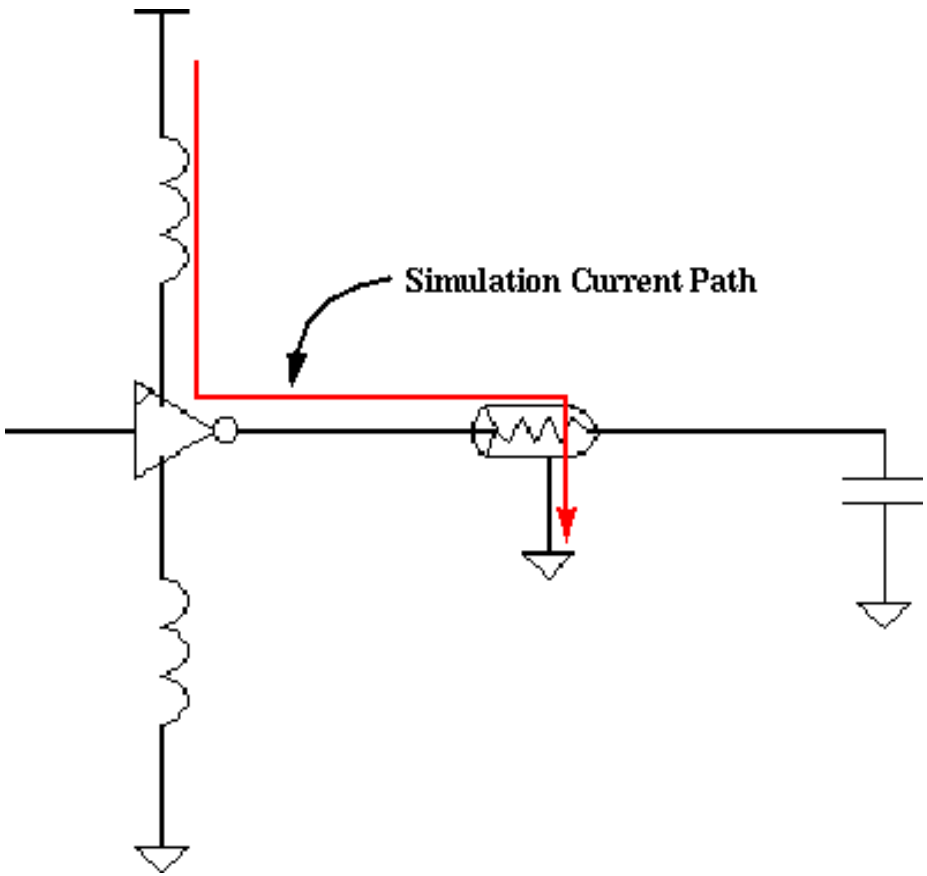

Filtered Supplies for Special Circuits

- Certain sensitive circuits need very quiet supplies to ensure faithful operation
  - Examples are PLL's and DLL's
- Filters can be used to control power noise
- A simple solution is an LC filter
  - Hooked up in a manner to convert differential noise to common model
  - Note that Vssa is *not* shorted to Vss

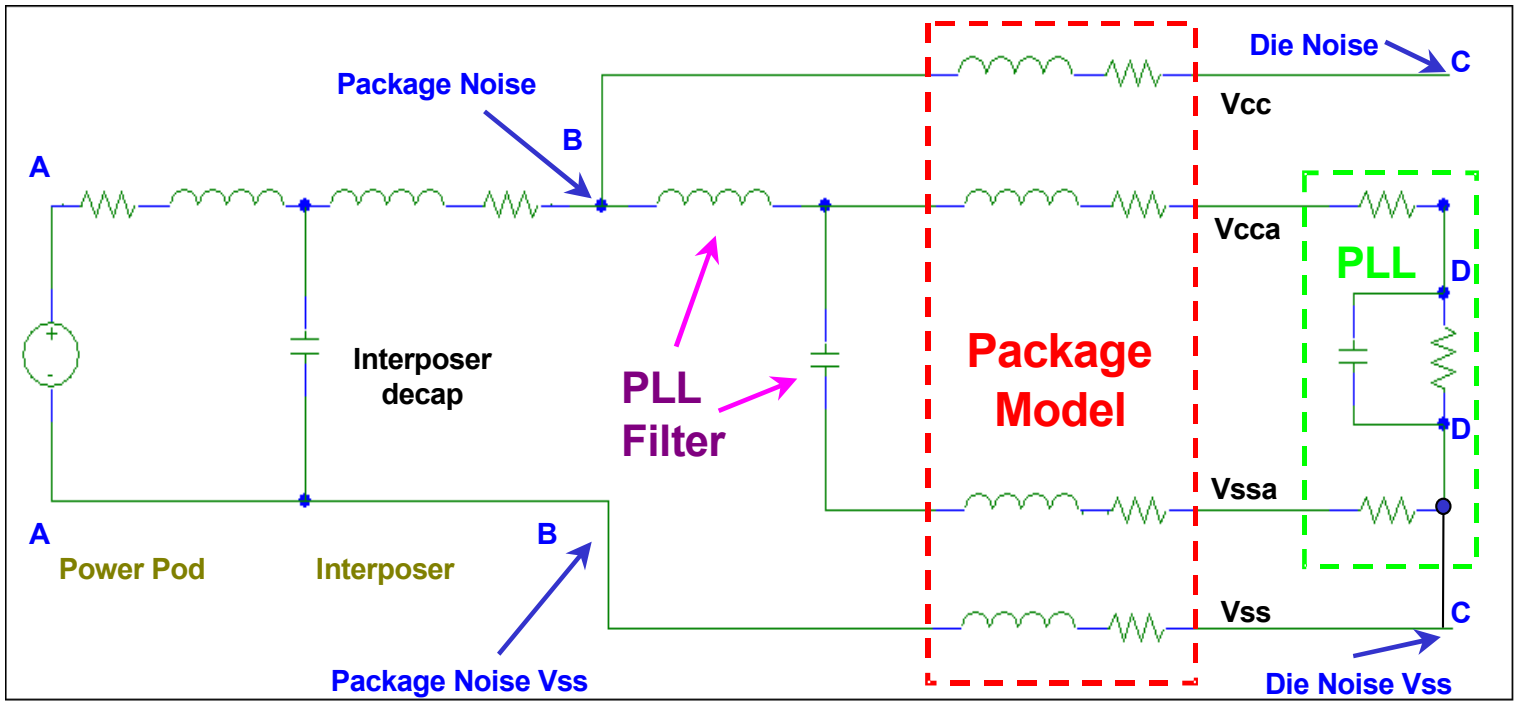

Die Power Delivery Scaling

- Transistor current scales on the time axis with channel length (actual, not drawn)
  - Current ramp will speed up by more than 0.7x per generation
- Capacitance density increases by about 1.4x per generation
- Wire resistance per unit length increases by ~ 2x per generation for scaled wires
  - Narrow wire width effects are making this worse
  - Retaining large dimensions for the top 2 layers (or adding layers) needed to offset this problem
    - Ensure good connectivity to top layer power
    - Otherwise RC delay for charge sharing will degrade
- RC per unit of length of power grid is fairly constant
- Unfortunately, RC path delay scaling is forcing more repeaters and increased sizes which exacerbates the problem
  - Bus width increases also compound the issue

Current Density Scaling

- Current density is C/Area*fV(AF)
- C/Area goes up ~ 1.4x
- Voltage has been trending down ~ 0.85x
  - Future decreases (beyond 90 nm) may be less (0.9 – 0.95x) due to leakage limitations
- Processor frequencies have been going up ~ 1.8x
  - Less for other types of IC's
  - Less in the future (~1.6x)?
- Leakage is increasing total current ~ 1.1x
- Activity factors ~ constant in worst areas
  - Decreasing on a full chip basis as cache area increases

- Overall current density increasing by ~ 2x
  - Forcing power metal needs to increase each generation

References

[1] Ng, et al., Table 1, paper 9.6, IEDM 2002

[2] Ishikawa, et al., paper 9.7, IEDM 2002

[3] Kedzierski, et al., paper 10.1, IEDM 2002

[4] Krivokapic, et al., paper 10.7, IEDM 2002

[5] Ang, et al., “An On-chip Voltage Regulator using Switched Decoupling Capacitors.” paper 26.7, ISSCC 2000

[6] Grochowski, E.; Ayers, D.; Tiwari, V. "Microarchitectural Simulation and Control of Di/Dt-Induced Power Supply Voltage Variation". *Proceedings of the Eighth International Symposium on High Performance Computer Architecture, 2002*. Page(s): 7-16.