Impact of Variability on Power

Sani R. Nassif
IBM Research – Austin
nassif@us.ibm.com

Realities

- Power has emerged as the #1 limiter of design performance beyond the 65nm generation.
- Dynamic and static power dissipation limit achievable performance due to fixed caps on chip or system cooling capacity.
- Power related signal integrity issues (IR drop, L di/dt noise) have become major sources of design re-spins.

ITRS Power Trends

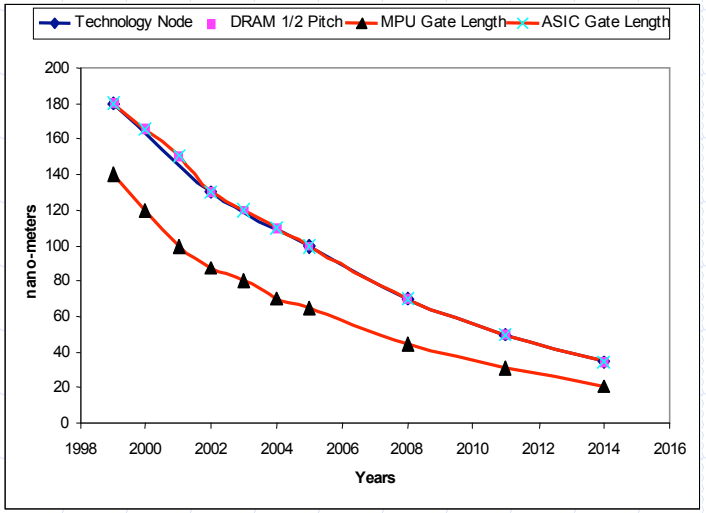

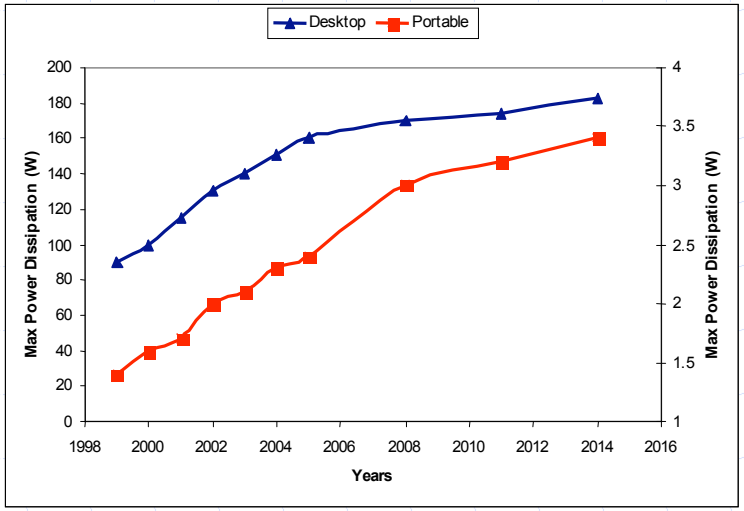

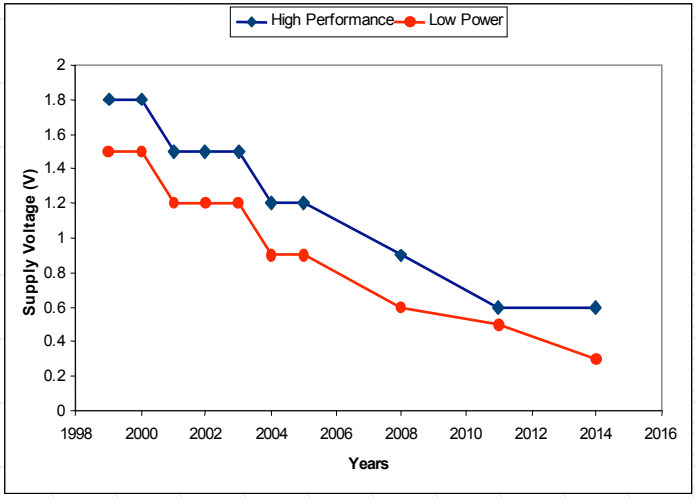

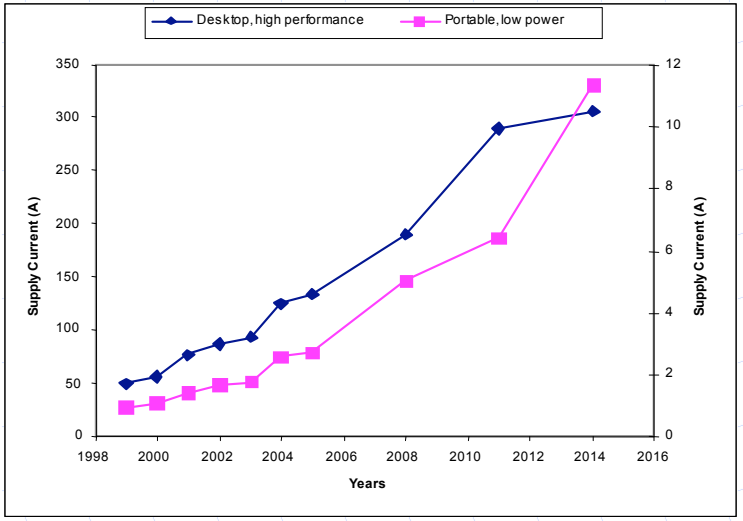

Farid Najm (from *TRS data)

Leakage Current "Predictions"

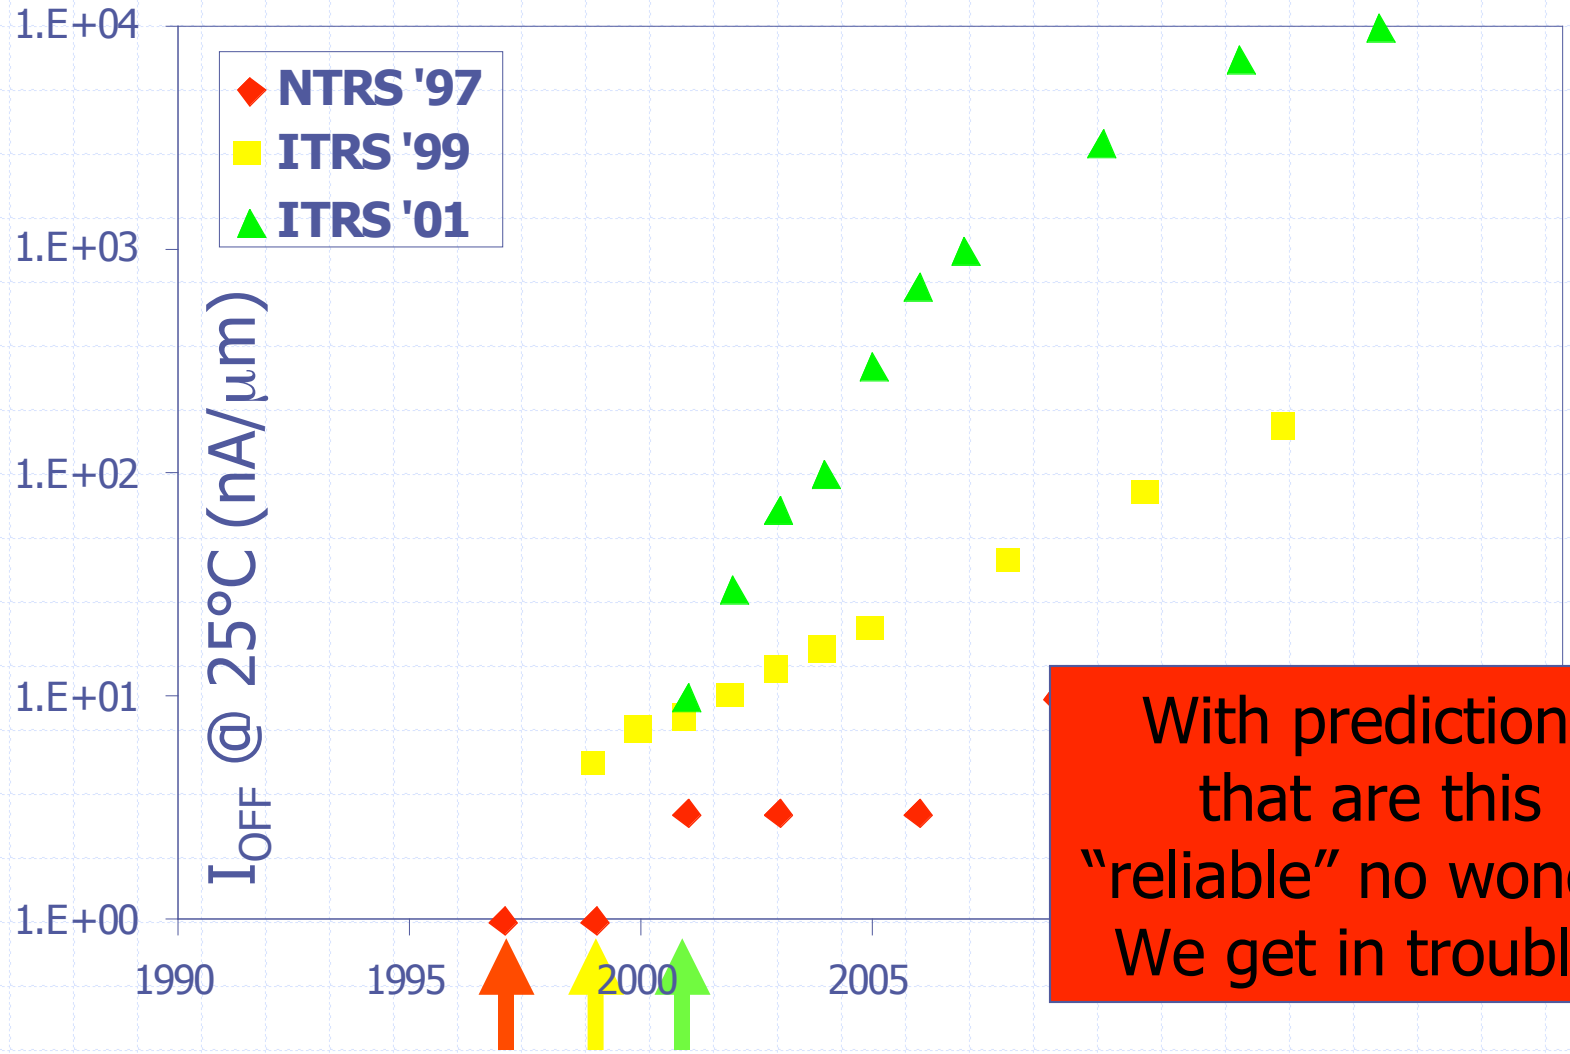

With predictions that are this "reliable" no wonder We get in trouble!

Tak Ning (from *TRS data)

Industry Views (Intel)

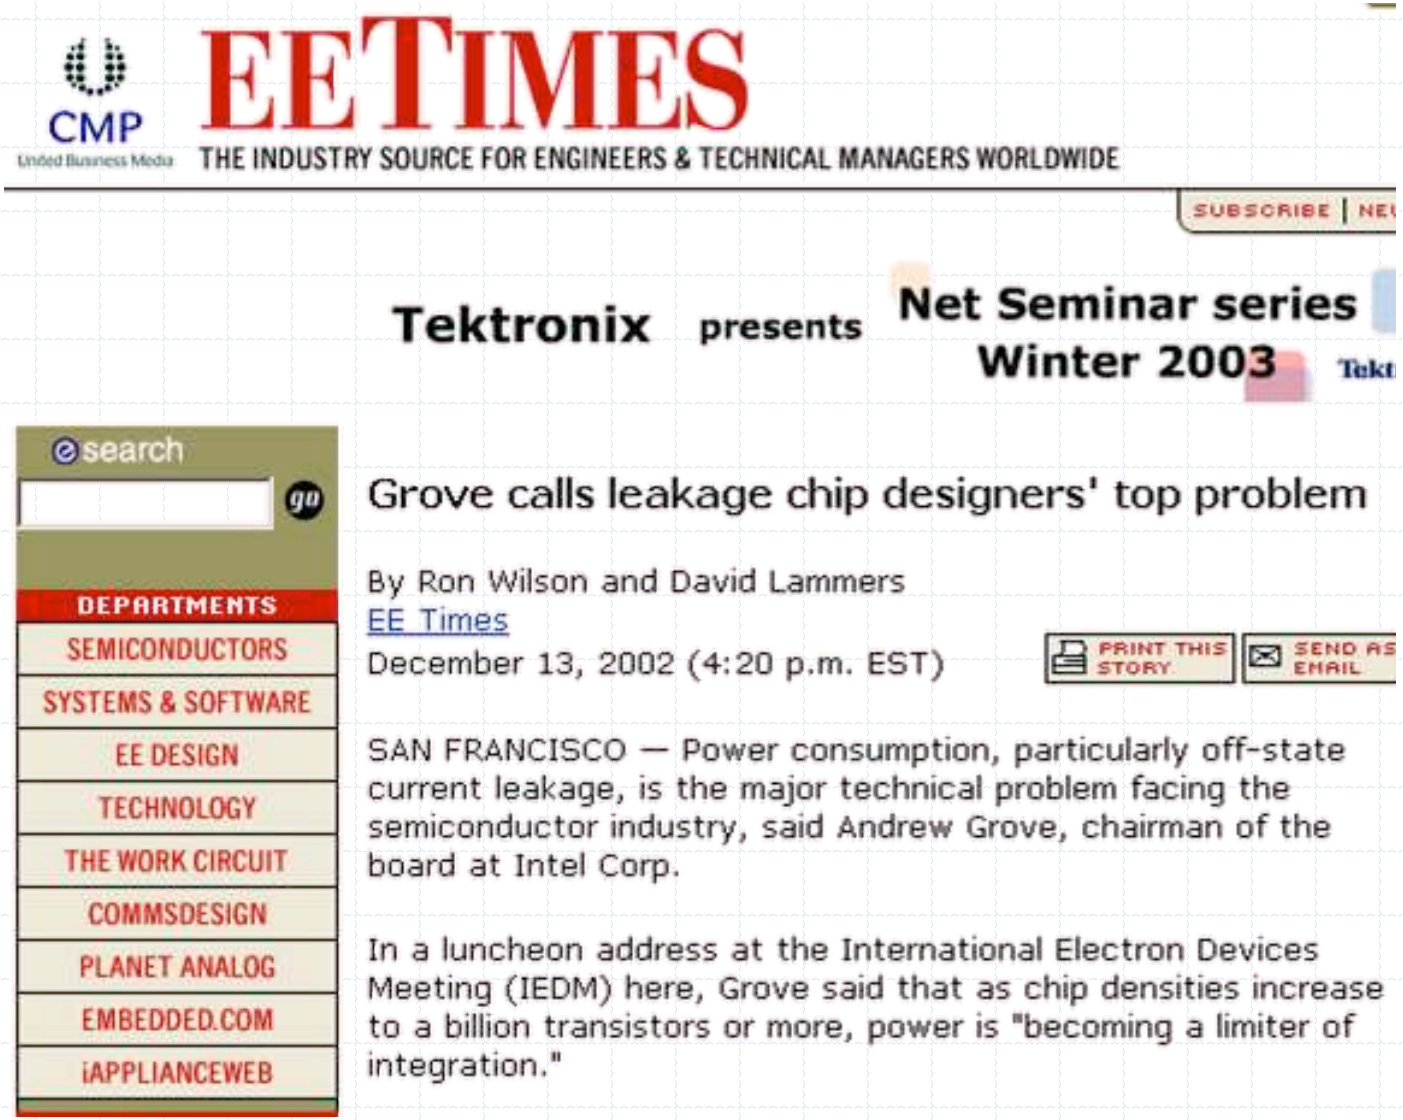

CMP
United Business Media

EETIMES
THE INDUSTRY SOURCE FOR ENGINEERS & TECHNICAL MANAGERS WORLDWIDE

SUBSCRIBE | NE

Tektronix presents Net Seminar series
Winter 2003 Tekt

search
go

DEPARTMENTS
SEMICONDUCTORS
SYSTEMS & SOFTWARE
EE DESIGN
TECHNOLOGY
THE WORK CIRCUIT
COMMSDESIGN
PLANET ANALOG
EMBEDDED.COM
iAPPLIANCEWEB

Grove calls leakage chip designers' top problem

By Ron Wilson and David Lammers
EE Times
December 13, 2002 (4:20 p.m. EST)

PRINT THIS STORY | SEND AS EMAIL

SAN FRANCISCO — Power consumption, particularly off-state current leakage, is the major technical problem facing the semiconductor industry, said Andrew Grove, chairman of the board at Intel Corp.

In a luncheon address at the International Electron Devices Meeting (IEDM) here, Grove said that as chip densities increase to a billion transistors or more, power is "becoming a limiter of integration."

Trends and Needs

- Technology trends:
  - Power and Frequency are increasing.
  - V_{DD} is decreasing (V_{TH} slower to manage leakage).
  - I_{DD} increasing (reliability/electromigration!).
- Impact of these trends:
  - IR and L di/dt have more impact on noise.
  - V_{DD} variation has more impact on delay.
- Critical need:
  - Understand supply induced noise variability and its future trends.

Acknowledgements

- IBM's ASIC power delivery team:
  - Raju Balasubramanian, Joseph Kozhaya, Bob Proctor, Erich Schanzenbach, Ivan Wemple.
- Michael Beattie, Byron Krauter, and Hui Zheng for the package variability analysis.
- Juan Antonio Carballo for power grid planning.

Outline

- Introduction
- Variability & Uncertainty
- Power Delivery Components
  - Circuits (power dissipation)
  - Decoupling Capacitance
  - On-Chip Power Grid
  - Package
- Tool Requirements
- Conclusions

Variability Sources

Physical

- Changes in characteristics of devices and wires.
- Caused by IC manufacturing process & wear-out (electro-migration).
- Time scale: 10^{9}sec (years).

Environmental

- Changes in V_{DD}, Temperature, local coupling.
- Caused by the specifics of the design implementation.
- Time scale: 10^{-6} to 10^{-9}sec (clock tick).

Variability Time Scales

$10^{-10} – 10^{-8}$	$10^{-7} – 10^{-5}$	$10^{-4} – 10^{-2}$	$10^{5} – 10^{7}$
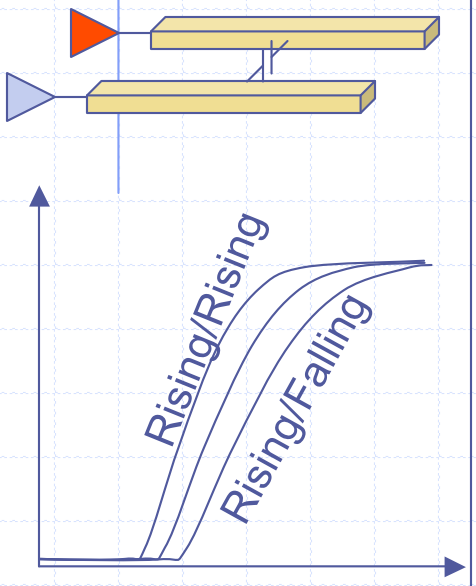	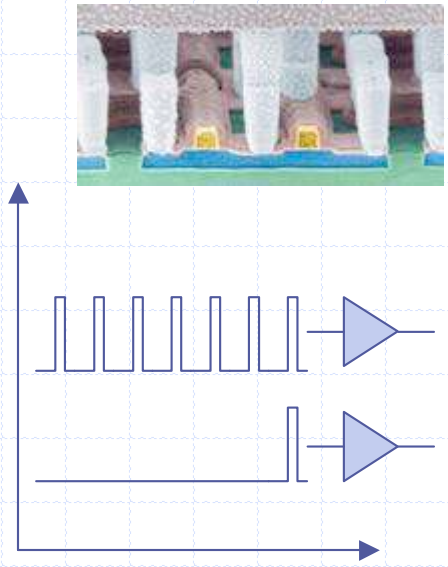	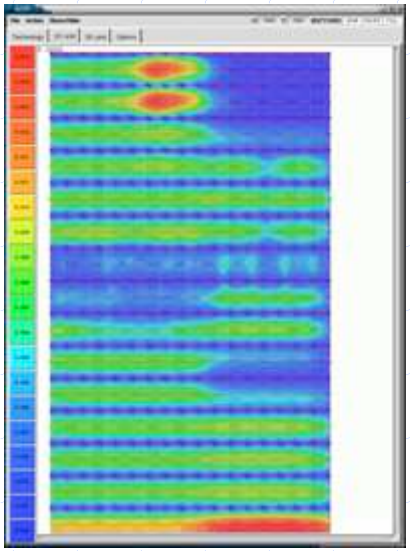	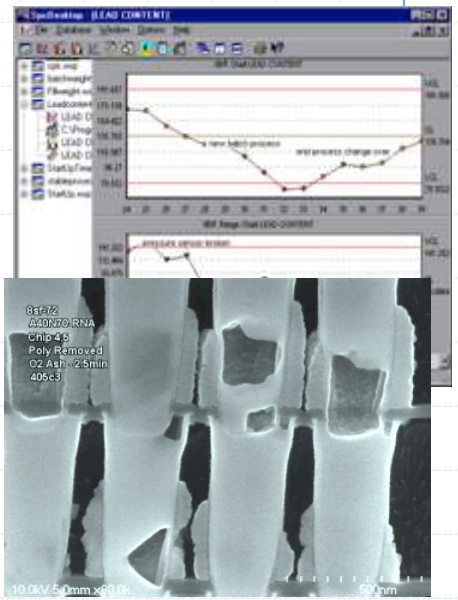
Signal Coupling IR drop	SOI History V_{DD}/Package Noise	Temperature	Process Line Electro-Migration

Variability Distribution

Physical

- Die to die variation
  - ***Imposed*** upon design (constant regardless of design).
  - Well modeled via worst-case files.
- Within-die variation
  - ***Co-generated*** between design & process (depend on details of the design).
  - Example: Via resistance variation vs. via density.

Environmental

- Only makes sense within-die.

Variability vs. Uncertainty

- Variability: known quantitative relationship to a source (readily modeled and simulated).
  - Designer has option to *null* out impact.
  - Example: power grid noise.
- Uncertainty: sources unknown, or model too difficult/costly to generate or simulate.
  - Usually treated by some type of worst-case analysis.
  - Example: ΔT_{OX} within die variation.
- Lack of modeling resources often transforms variability to uncertainty.
  - Example: switching probability assessment.

Uncertainty in Design-Process

- Design uncertainty:
  - Portions not yet defined.
  - Changes in specification.
- Modeling uncertainty:
  - Lack of detail in models.
  - Pessimism/conservatism.
- Processing uncertainty:
  - Manufacturing noise (ΔL, V_T).
  - Changes as technology matures.
  - Accuracy needed relatively late in the design cycle.

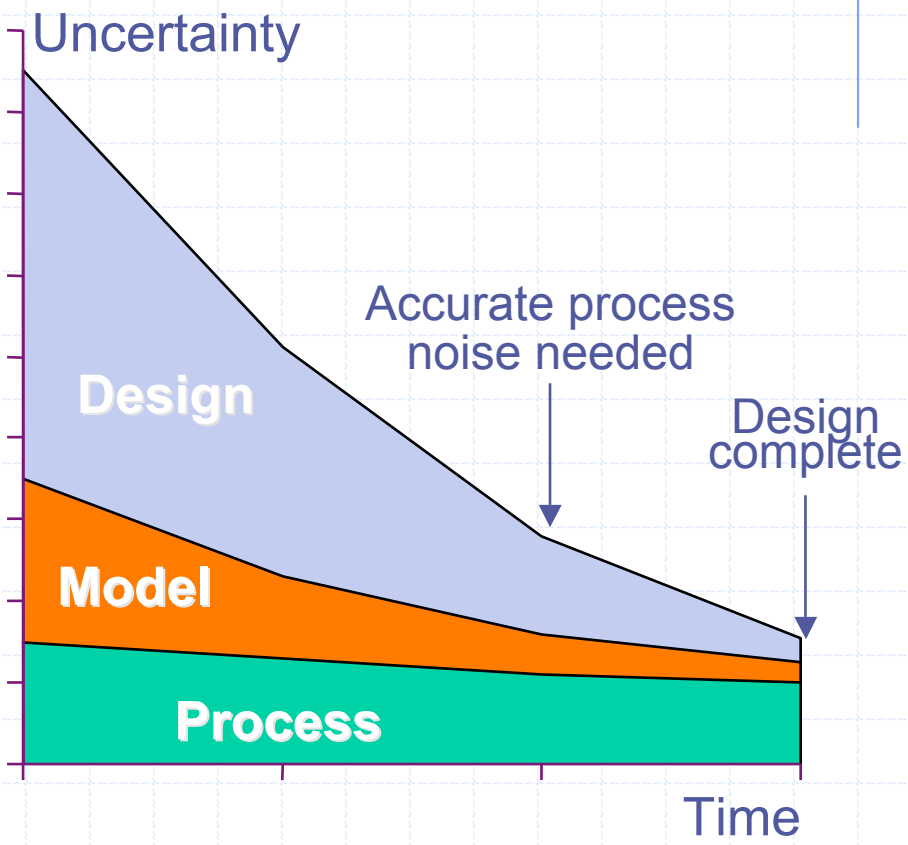

Variability & Uncertainty

- In the power delivery area, large amounts of uncertainty exist (more than for timing…).
- Circuit activity is seldom known well enough to allow accurate prediction.
  - Relatively well known fact.
- Little is known on the dependence of the various components of power on technology and its variability (hence this tutorial).

Outline

- Introduction
- Variability & Uncertainty
- Power Delivery Components
  - Circuits (power dissipation)
  - Decoupling Capacitance
  - On-Chip Power Grid
  - Package
- Tool Requirements
- Conclusions

Power Delivery Components

Package

On-Chip Power Grid

Circuits

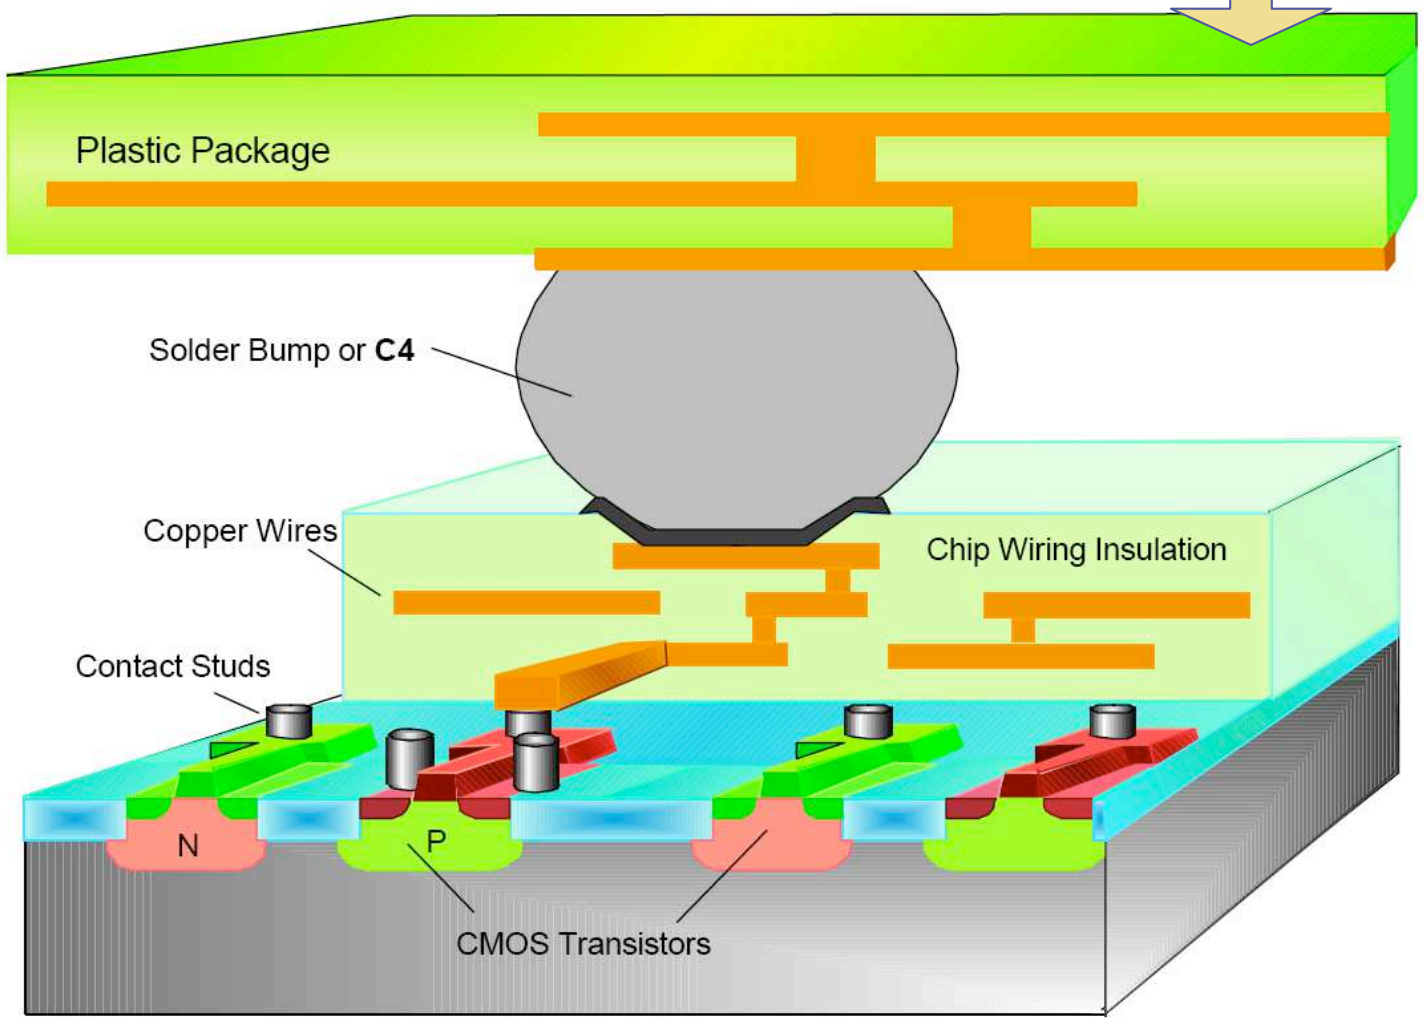

Power Delivery Components

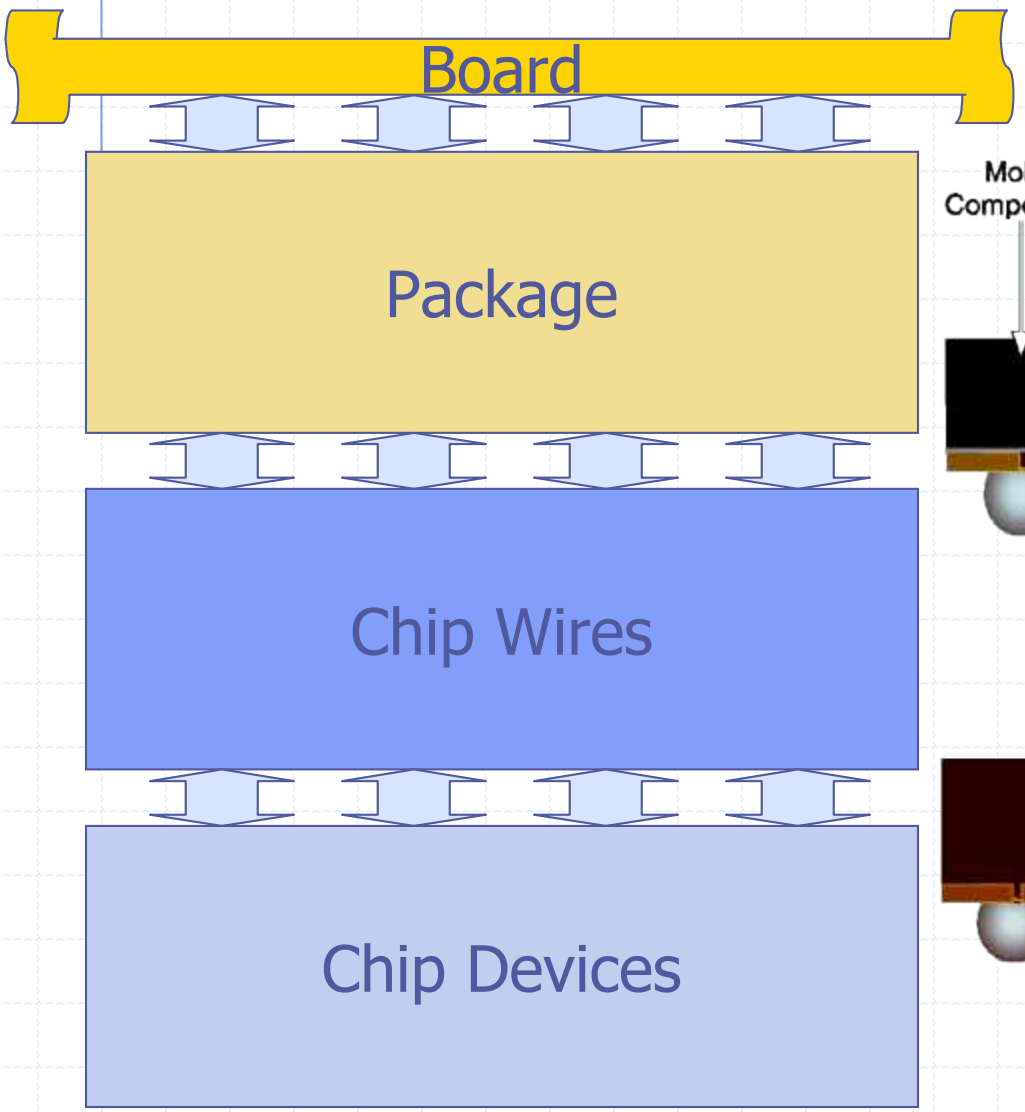

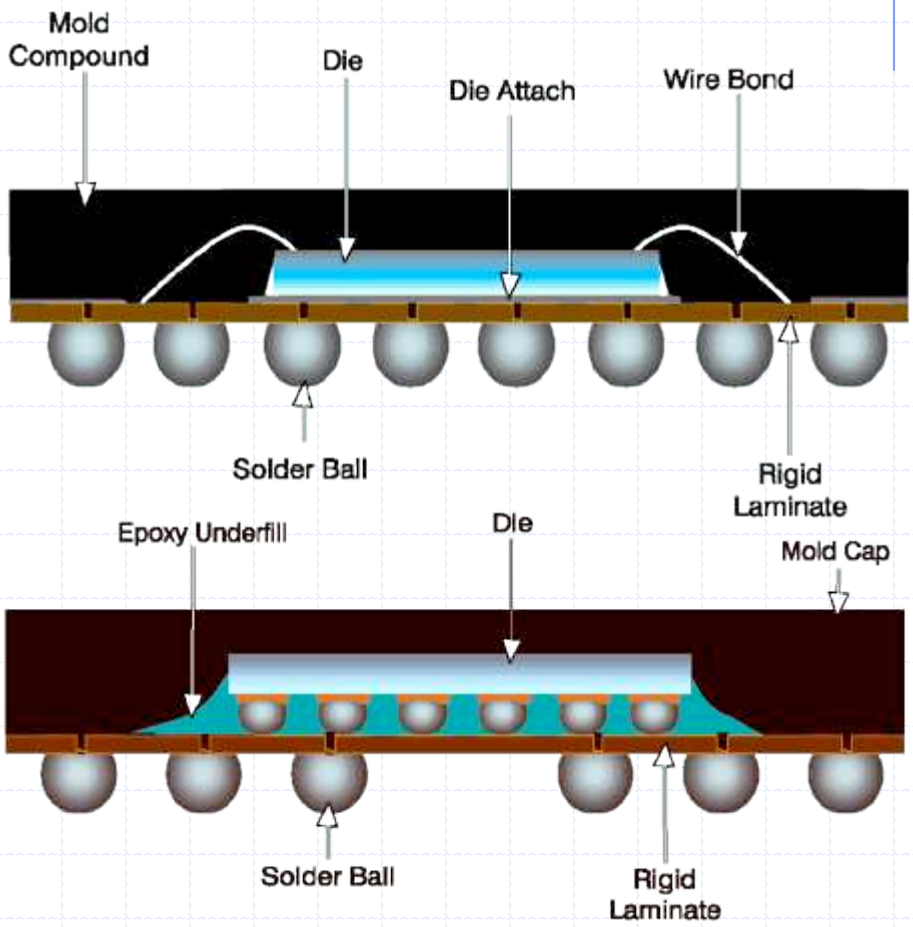

http://www.lithium.it (via Google)

Power Variability Components

- Board level, not addressed in this work.
- Package level.
  - Variability in package parasitics (R & L).
- On-Chip Power Grid level.
  - Variability in grid parasitics (R).
- Circuit level.
  - Variability in static and dynamic power consumed.
  - Variability in decoupling capacitance.

Example: Power Grid Noise

- Grid is predominantly **resistive**.
- Package is predominantly **inductive**.
- Load is modeled as a current.
- Other circuits ~ lossy decoupling capacitance.

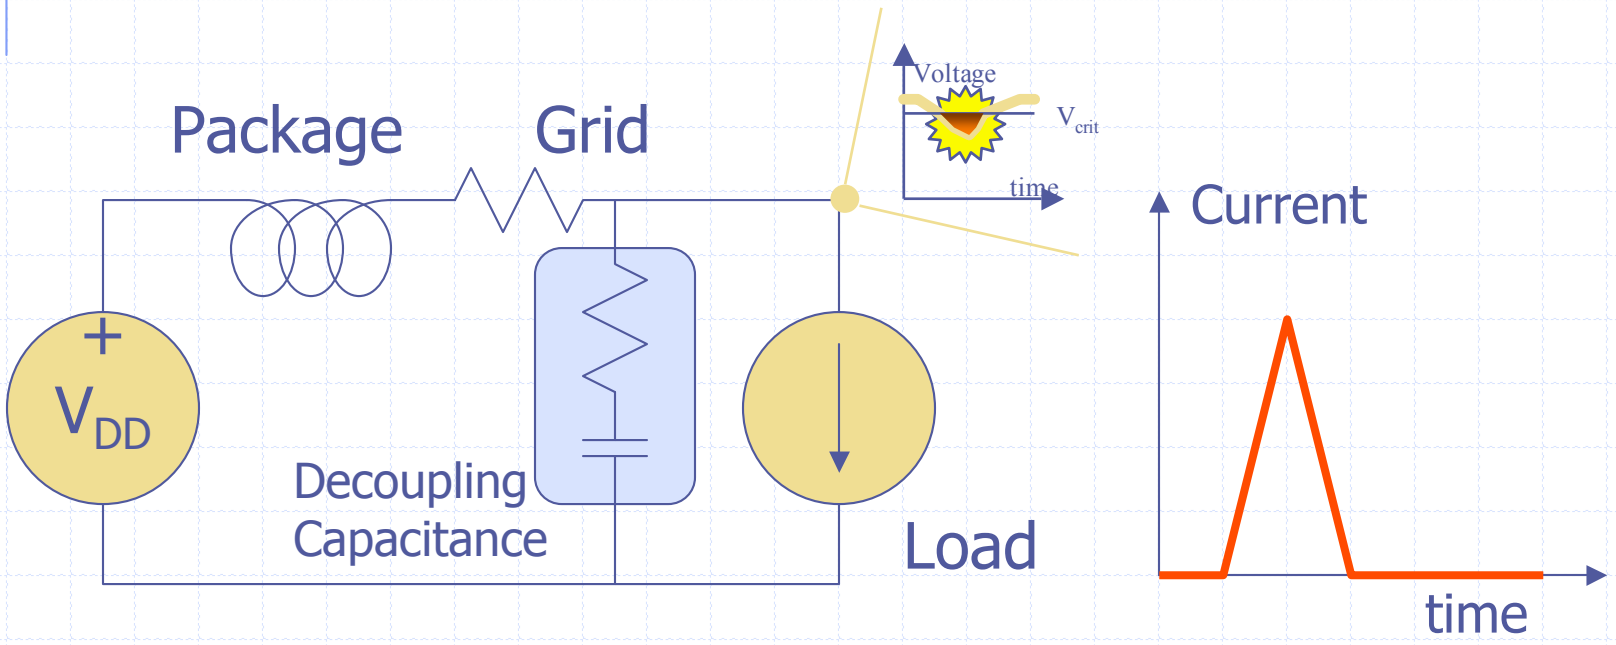

Noise Model

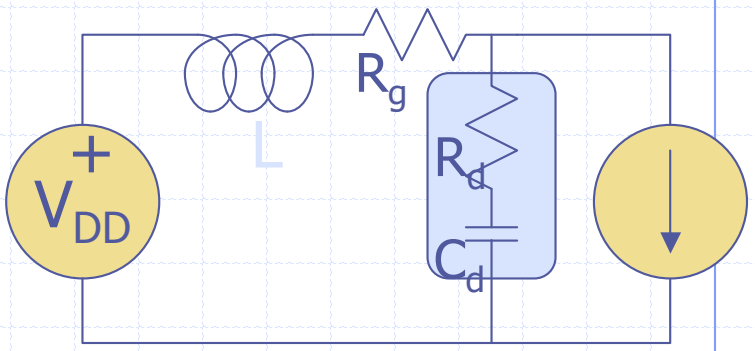

- Current modeled as:

$I = 0$	$t < 0$
$I = \mu t$	$t < t_p$
$I = \mu(2t_p - t)$	$t < 2\,t_p$
$I = 0$	$t > 2\,t_p$

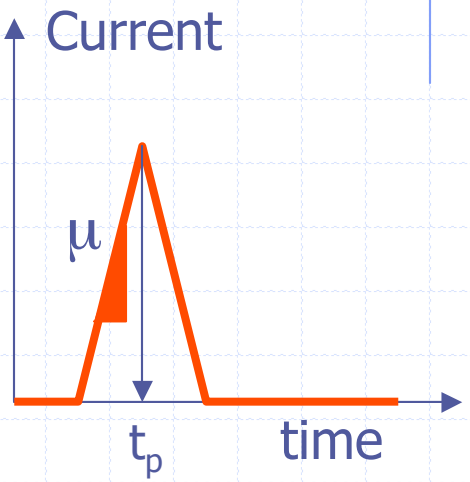

- Ignoring L, maximum noise is:

$$V_{max} = \underbrace{\mu\, t_p\, R_g}_{DC} - \underbrace{\mu\, R^2_g\, C_d\,(1 - e^{-t_p/\tau})}_{Decap}$$

$\tau = (R_g + R_d)\, C_d$

$\lambda\, \mu\, t_p\, R_g / (R_g + R_d)$
(for large C_d)

Noise Model + L

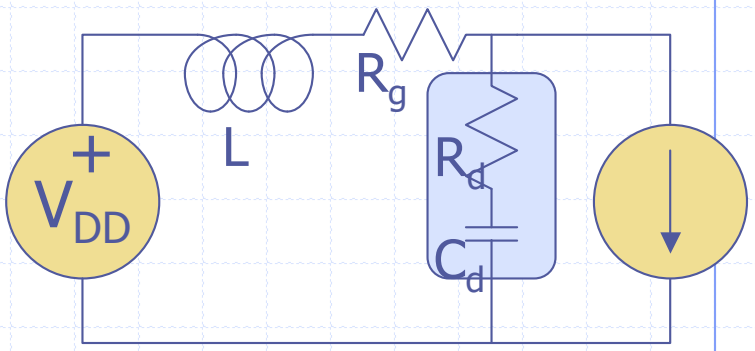

- With package, maximum noise becomes:

$$V_{max}\,\lambda\,\underbrace{\mu\,t_p\,R_g}_{DC}+\underbrace{\mu L}_{Package}-\underbrace{\mu\,R^2_g\,C_d\,(1-e^{-t_p/\tau})}_{Decap}$$

- Accurate expression:

$$V_{max}=\mu\,t_p\,R_g+\mu L-\mu\,R^2_g\,C_d+\psi_1+\psi_2$$

$$e_1=\exp -(\tau+\beta)t_p/2C_dL \qquad e_2=\exp -(\tau-\beta)t_p/2C_dL$$

$$\beta=(\tau^2-4LC_d)^{1/2}$$

$$\psi_1=(e_1+e_2)\,\mu\,(L-C_d\,R_g^2)/2$$

$$\psi_2=(e_1-e_2)\,\mu\,C_d\,(\tau\,R_g^2-L(3R_g-R_d))/2\beta$$

Noise Trends

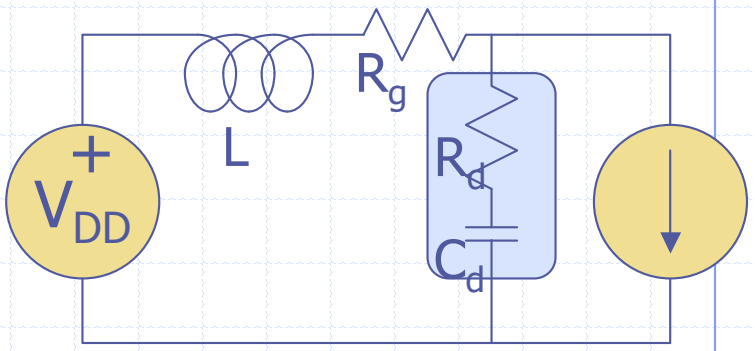

V_{DD} 0.6X

t_p 0.6X

μ 3.3X

C_d 2X

Based on conservative ITRS trends

~2X ~3X

$$V_{max} \lambda \mu t_p R_g + \mu L - \mu R^2_g C_d (1 - e^{-t_p/\tau}) \sim \text{Same}$$

DC *Package* *Decap*

Each of these components has a variability/tolerance!

Some Source of Power Variability

Package R/L vary with design & manufacturing tolerances.

Grid resistance varies due to CMP and other physical factors.

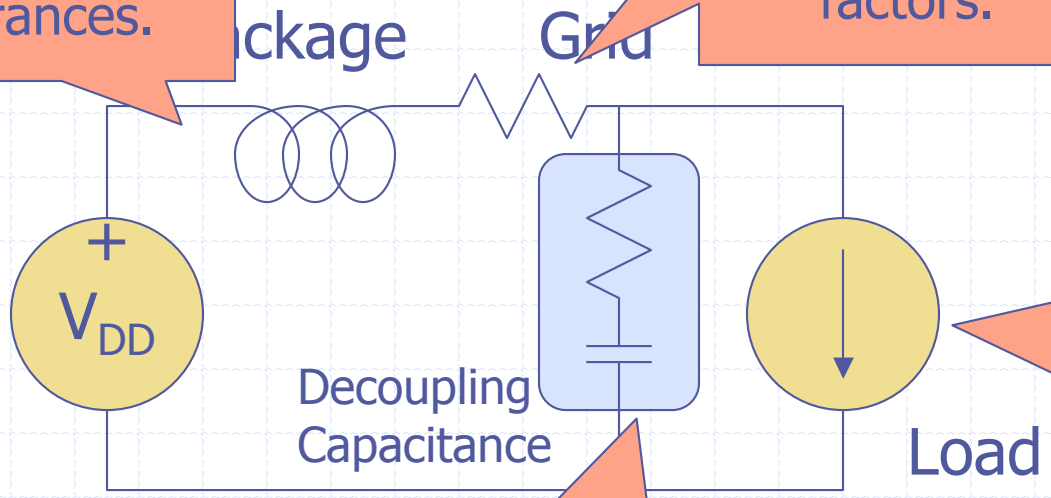

Load varies with operation, ΔL, V_T, and other physical factors.

Decap varies with operation, V_T, T_{OX}, and other physical factors

Outline

- Introduction
- Variability & Uncertainty
- Power Delivery Components
  - Circuits (power dissipation)
  - Decoupling Capacitance
  - On-Chip Power Grid
  - Package
- Tool Requirements
- Conclusions

Circuit Power Dissipation

- Circuits dissipate power when performing function.
  - $P \sim \alpha\ f\ C\ V^2$
  - C is a combination of wire and device capacitance (different sources of variability!).
  - V is the power supply (more than one may exist on a chip).
  - f is the frequency of operation.
  - The α factor models switching frequency & second order effects (e.g. short circuit current).

Circuit Power Breakdown

- For a modern digital system, power can be broken down by major type of component.
  1. Clock distribution and latches.
  2. Data-Path and custom logic.
  3. Arrays and Memories.
  4. Random Logic (cell based).
  5. Embedded cores (recursive).

- Each design has a different breakdown, so generalizations are not very useful.

Component Power Characteristics

- Each type of component lends itself to a specific form of power estimation, and hence power variability estimation.

- Example: Clocks.
  - Large amount of power (C ↑).
  - Highest frequency in the chip (f ↑).
  - Highest switching probability (α ↑).
  - Variability from technology, V_{DD} and from design specifics (clock gating).

Component Power Variability

- Rough breakdown of sensitivity to variability by component (mostly based on my opinion).
  - Denote highest 2 components… (YMWV).

	α	C wire	C device	V	f
Clock		X		X	
Data Path	X		X		
Arrays		X	X		
Random Logic	X		X		

Specifics: Random Logic

- Probably the hardest of the various components to analyze.
  - Inputs usually ill specified, requiring higher level architectural simulation in order to properly assess.

- BUT… usually a modest portion of overall chip power (for μ-Processor like designs).
  - Not often worth the effort except for designs with a large synthesized portion.

Example: ISCAS C432 Circuit

- Small combinational-only circuit.
  - 37 inputs, 7 output.
  - 160 gates, (1 AND 8, 3 AND 9, 64 NAND 2, 1 NAND 3, 14 NAND 4, 19 NOR 2, 40 NOT, 18 XOR 2).
- Strategy: explore process/pattern dependence.
  - Could have also looked at V_{DD} and T!
- Pattern: apply 50 random patterns (@1GHz) with 20% of the bits changing each cycle.
- Process: apply 58 unique sets of MOSIS 0.18μ CMOS parameters (represent wafer averages).

Example: I_{DD} Waveform

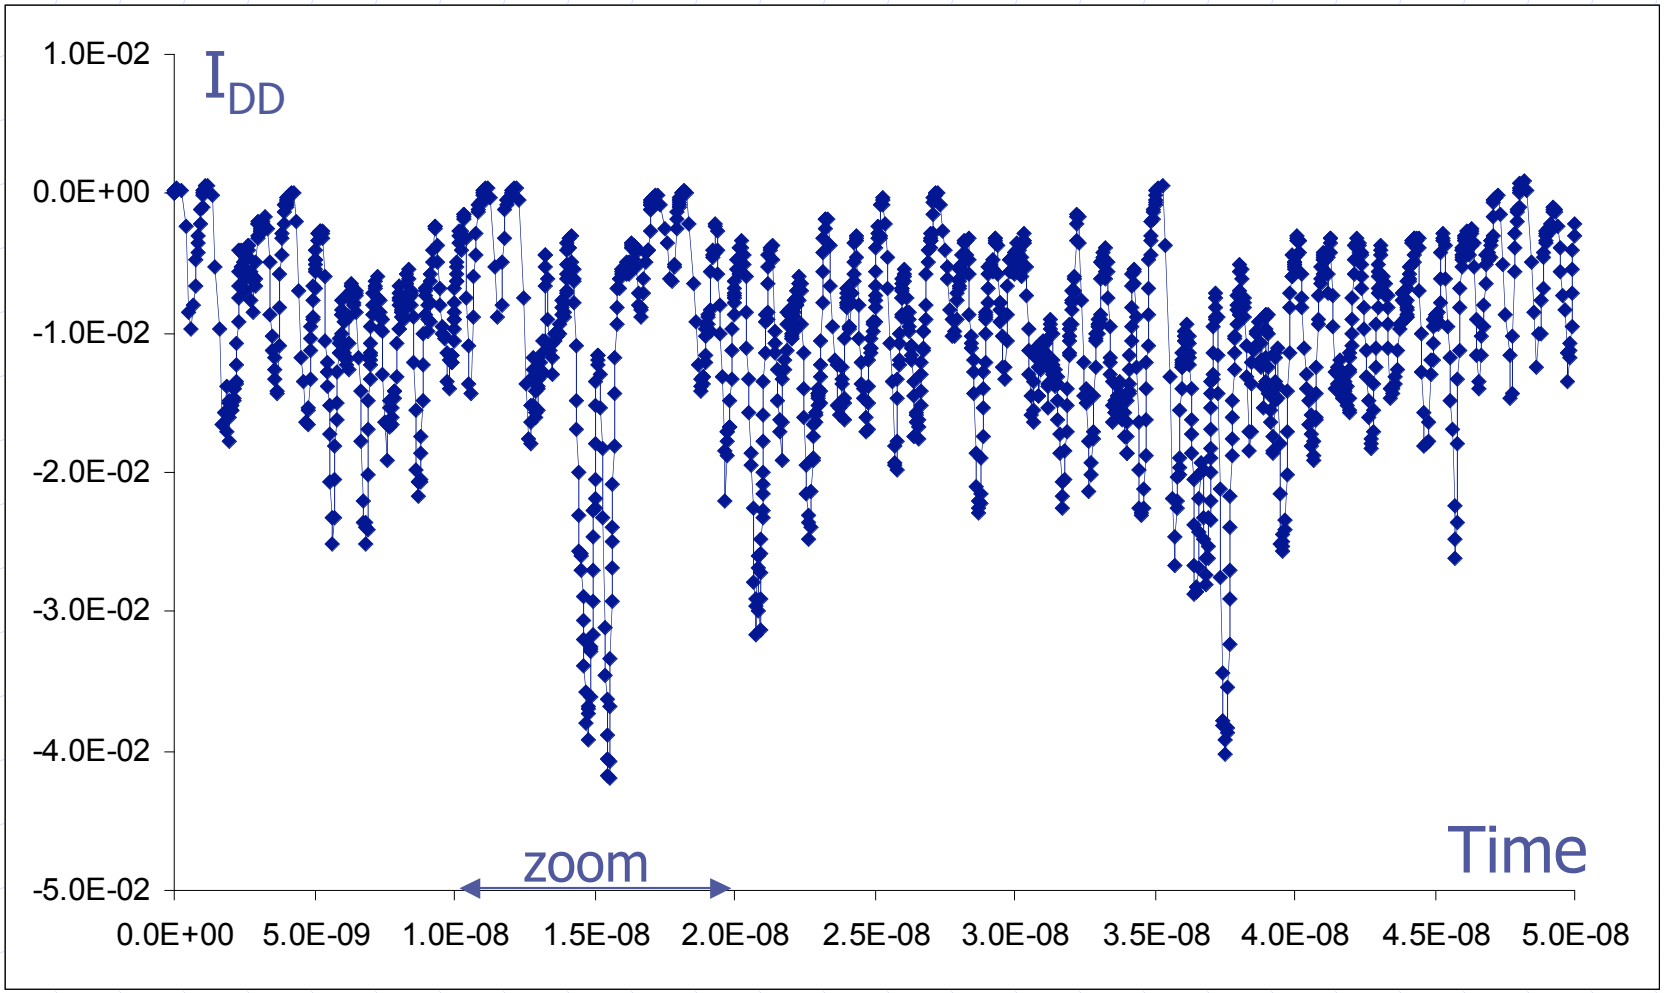

Example: Detail Waveform

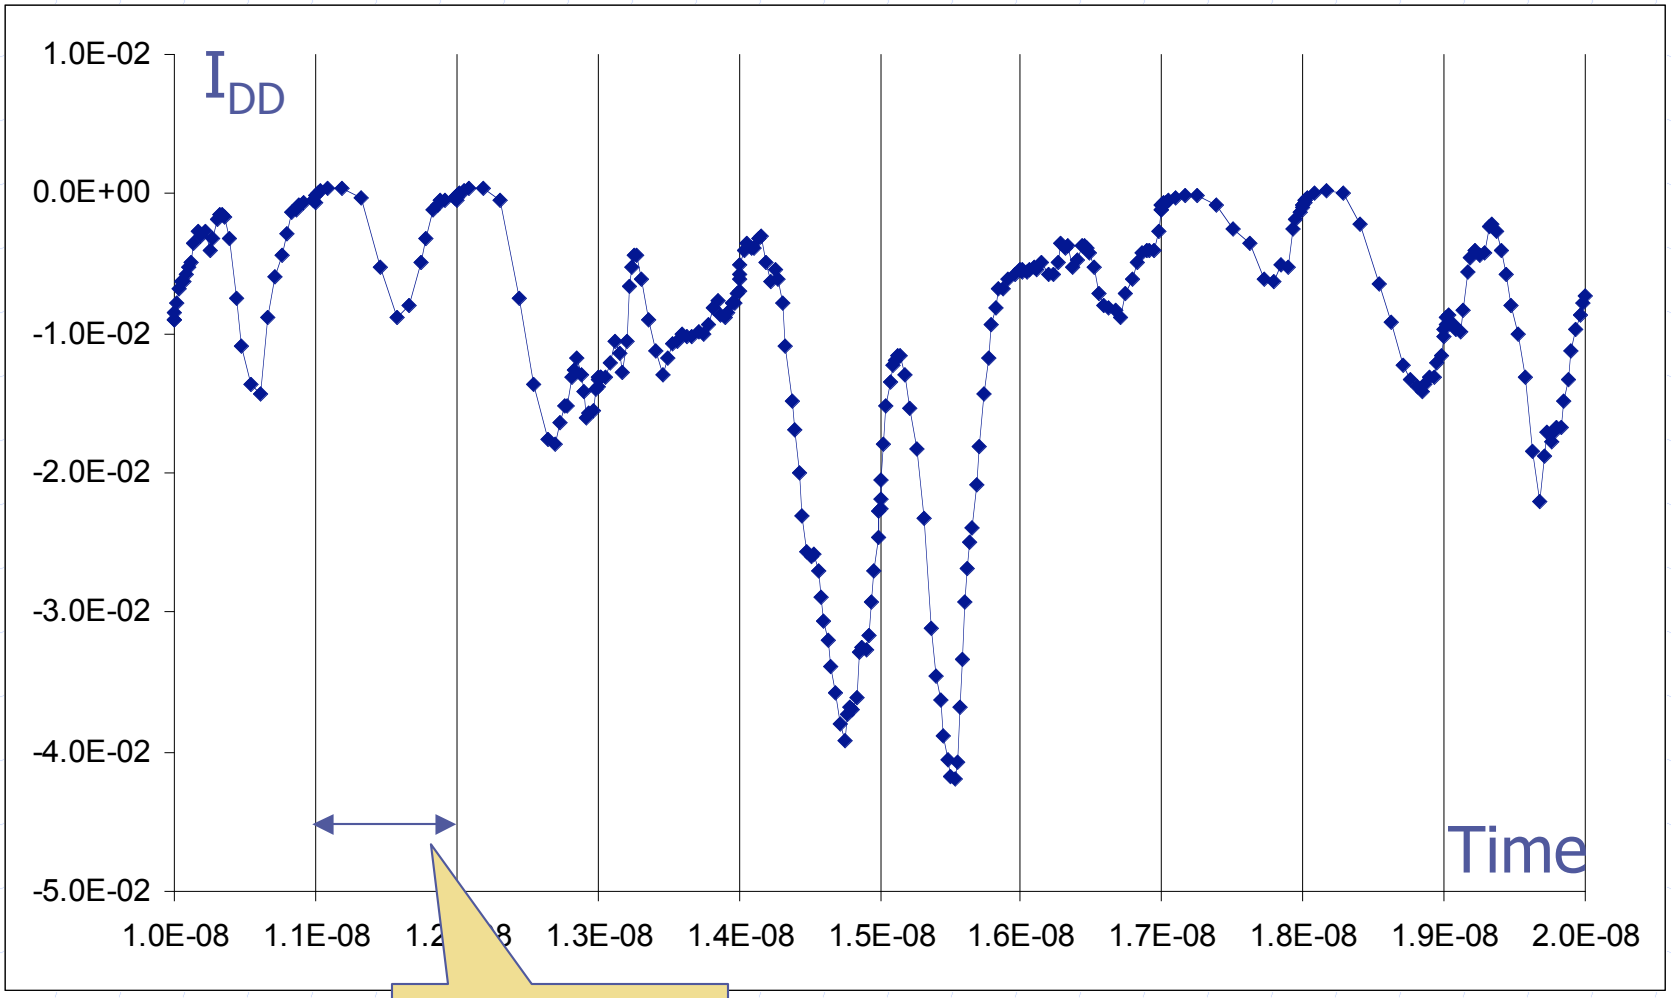

Pattern Influence Metric

- Compare peaks within each clock cycle.
  - Highest vs. lowest peak ∀ clock cycles.

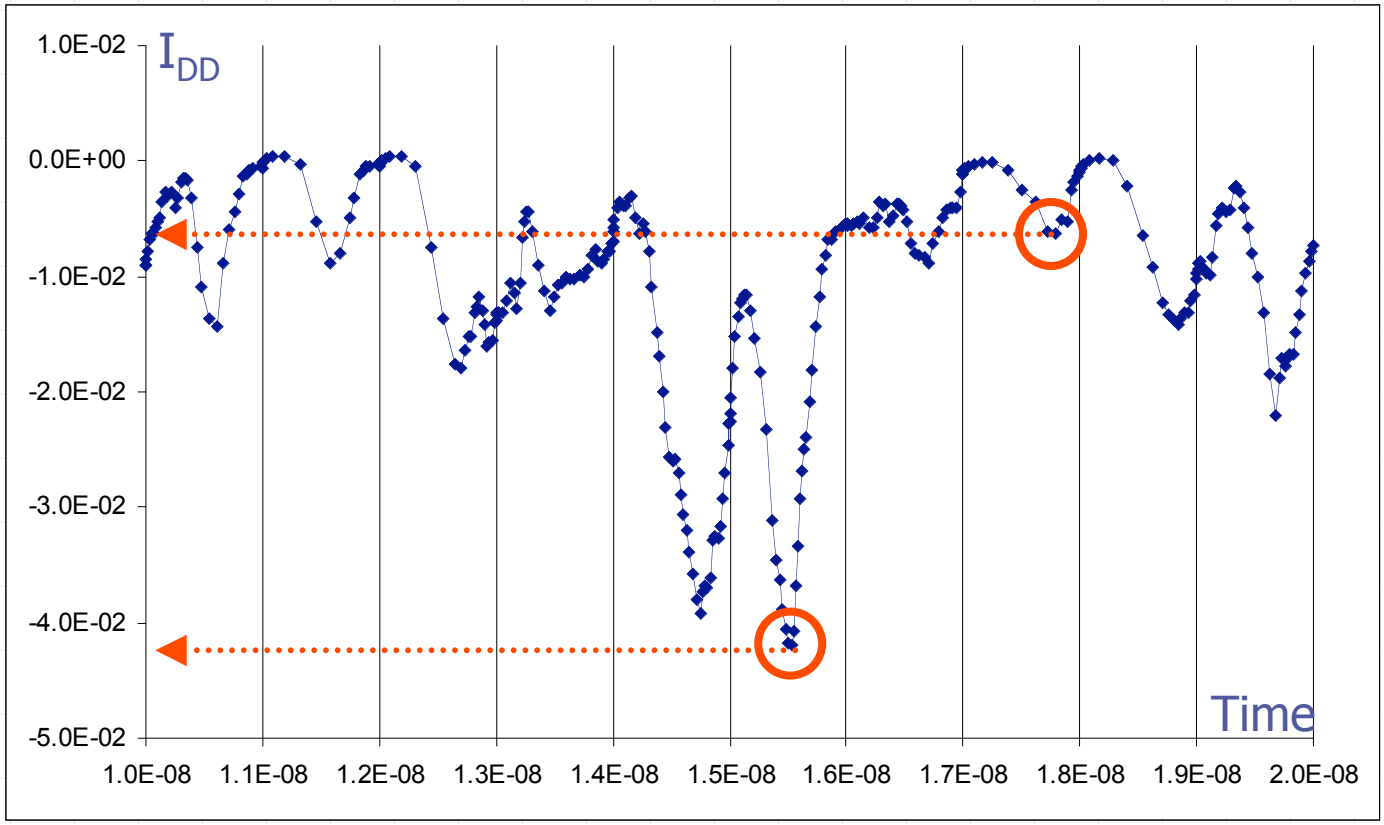

I_{DD} Peaks Per Clock Cycle

- Maximum absolute value of I_{DD} ∀ clock cycle.
- Max/Min range indicates overall influence.

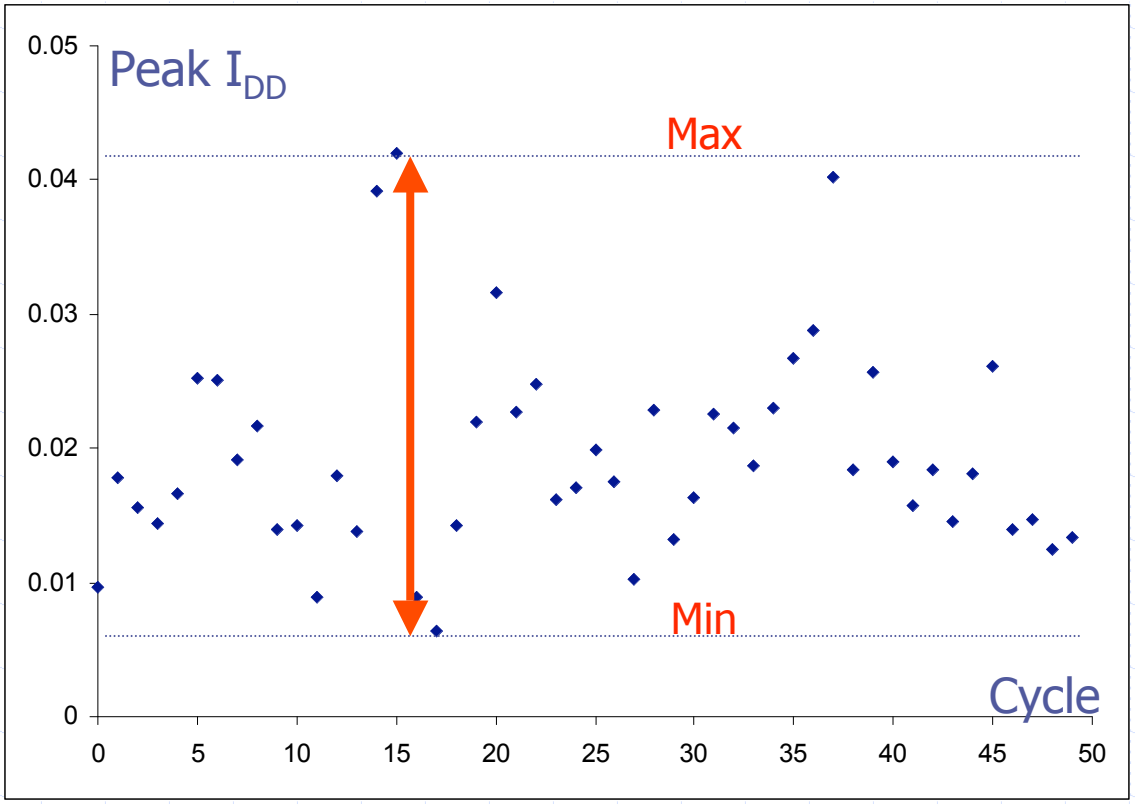

I_{DD} Peak Range Analysis

- Compare max-min range vs. process & pattern.

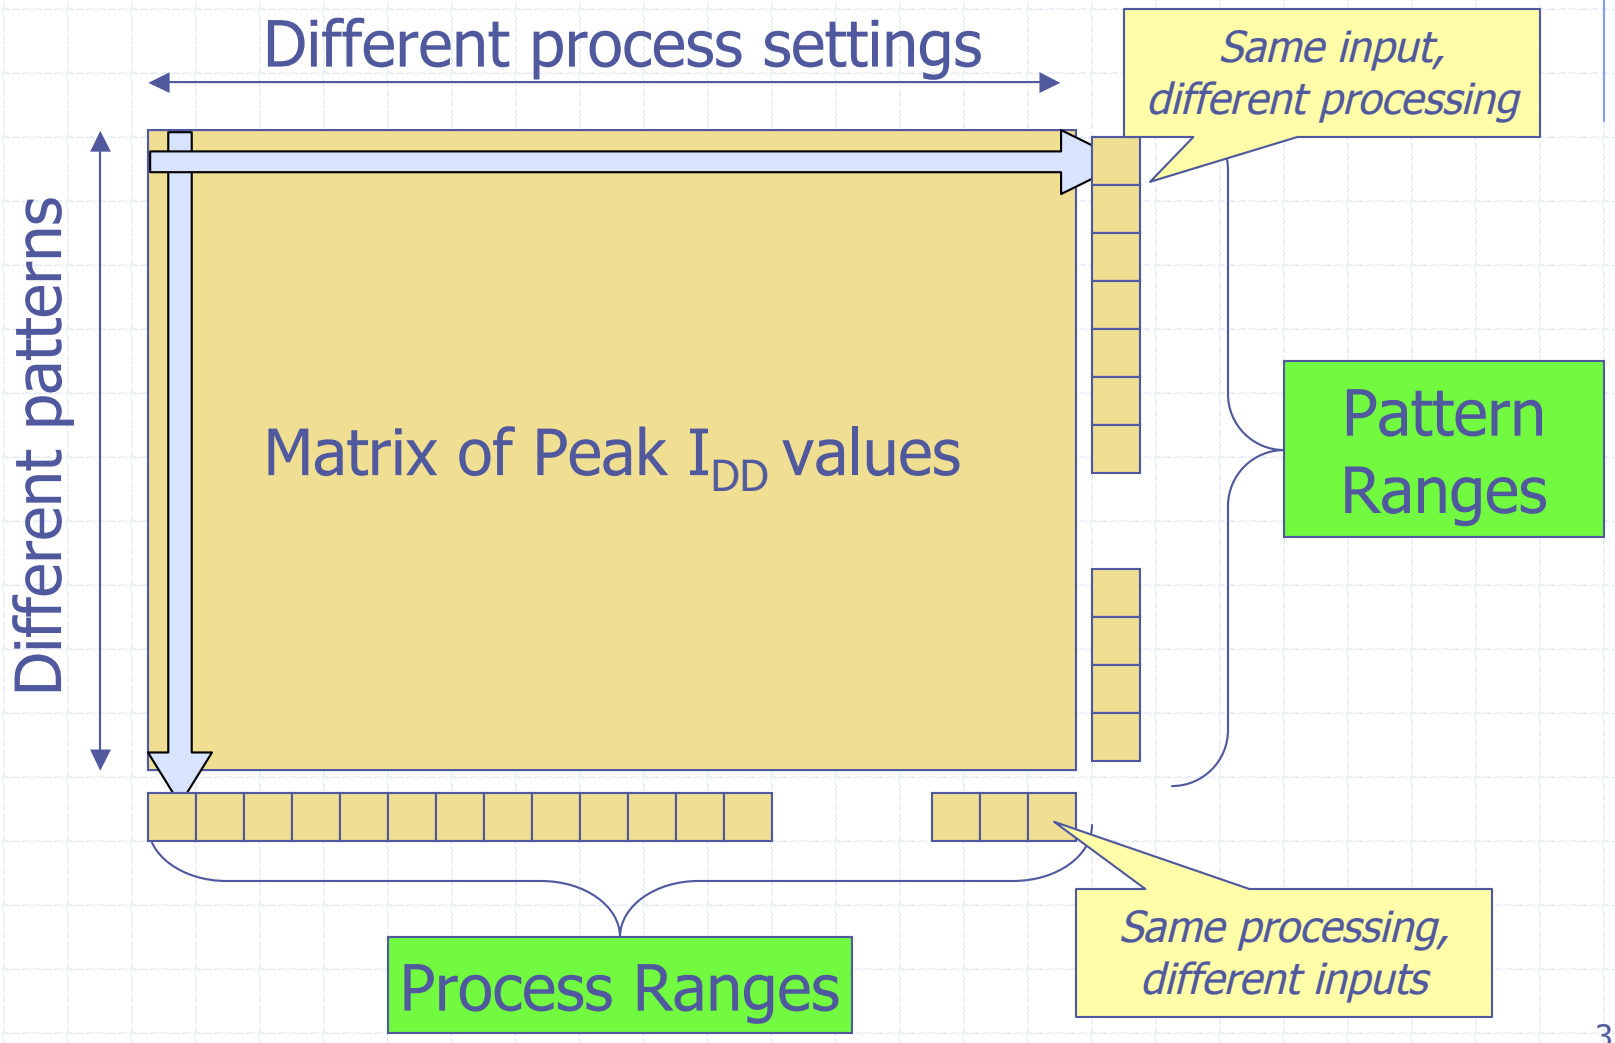

IDD Range Statistics

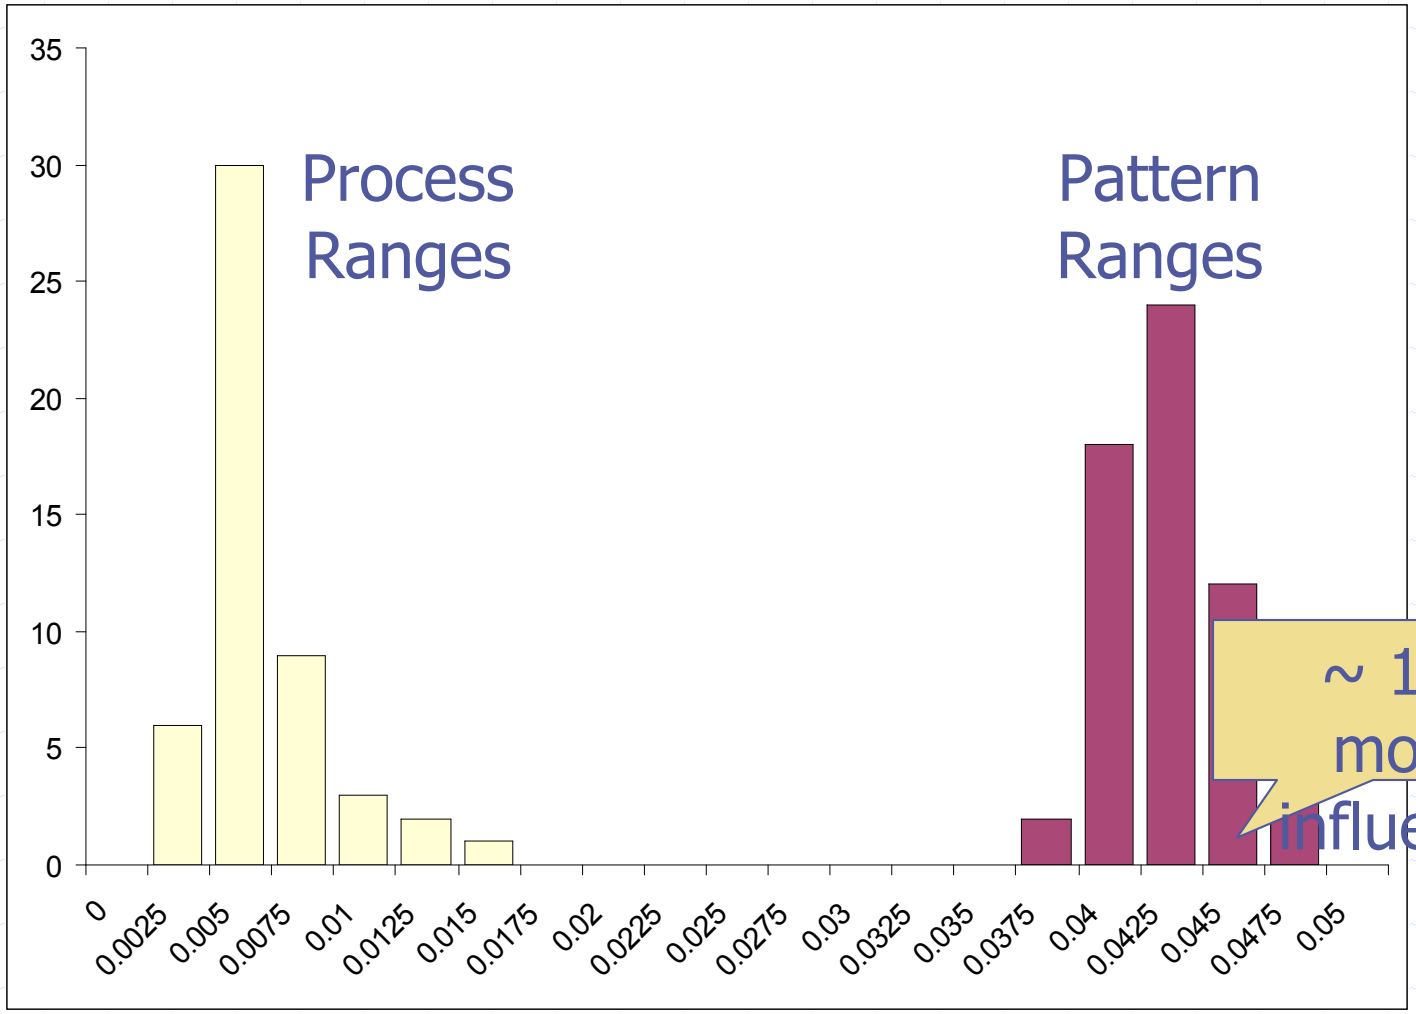

Conclusion

- For random logic, switching pattern dominates the variability.
  - Knowing the activity factor (α) is crucial.
- Since the process dependency is weak, higher level simulation (i.e. not Spice) can do well.
  - Details of process dependence can be abstracted away without too much loss of information.
- We have not looked at V_{DD} and Temperature!

More Conclusions

- For other components in the design, special purpose estimation techniques can be used to reduce the task to a few representative circuit simulations.
  - Example: cross-sections of a RAM + overall expected activity factors.

Outline

- Introduction
- Variability & Uncertainty
- Power Delivery Components
  - Circuits (power dissipation)
  - Decoupling Capacitance
  - On-Chip Power Grid
  - Package
- Tool Requirements
- Conclusions

Decoupling Capacitance

- MOSFETs have intrinsic and extrinsic linear and non-linear capacitances.
- When a circuit is not active, these capacitors act as a reservoir of charge which can be supplied to neighboring active circuits.

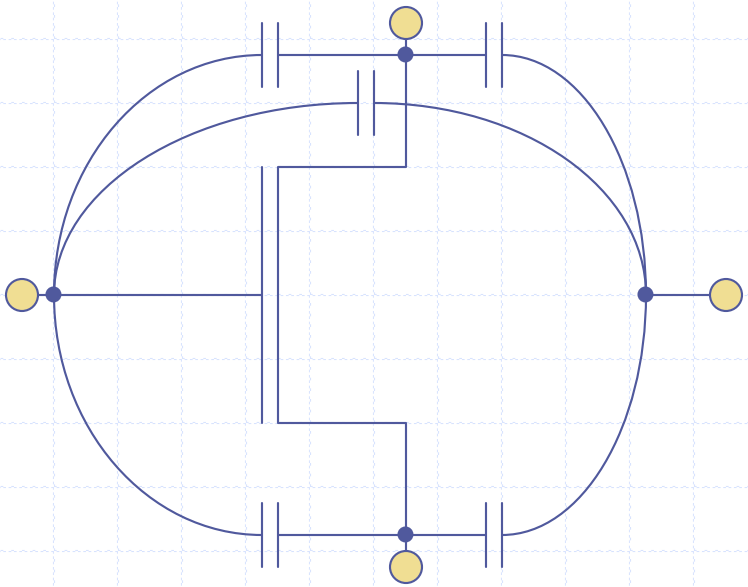

Decoupling Example: Inverter

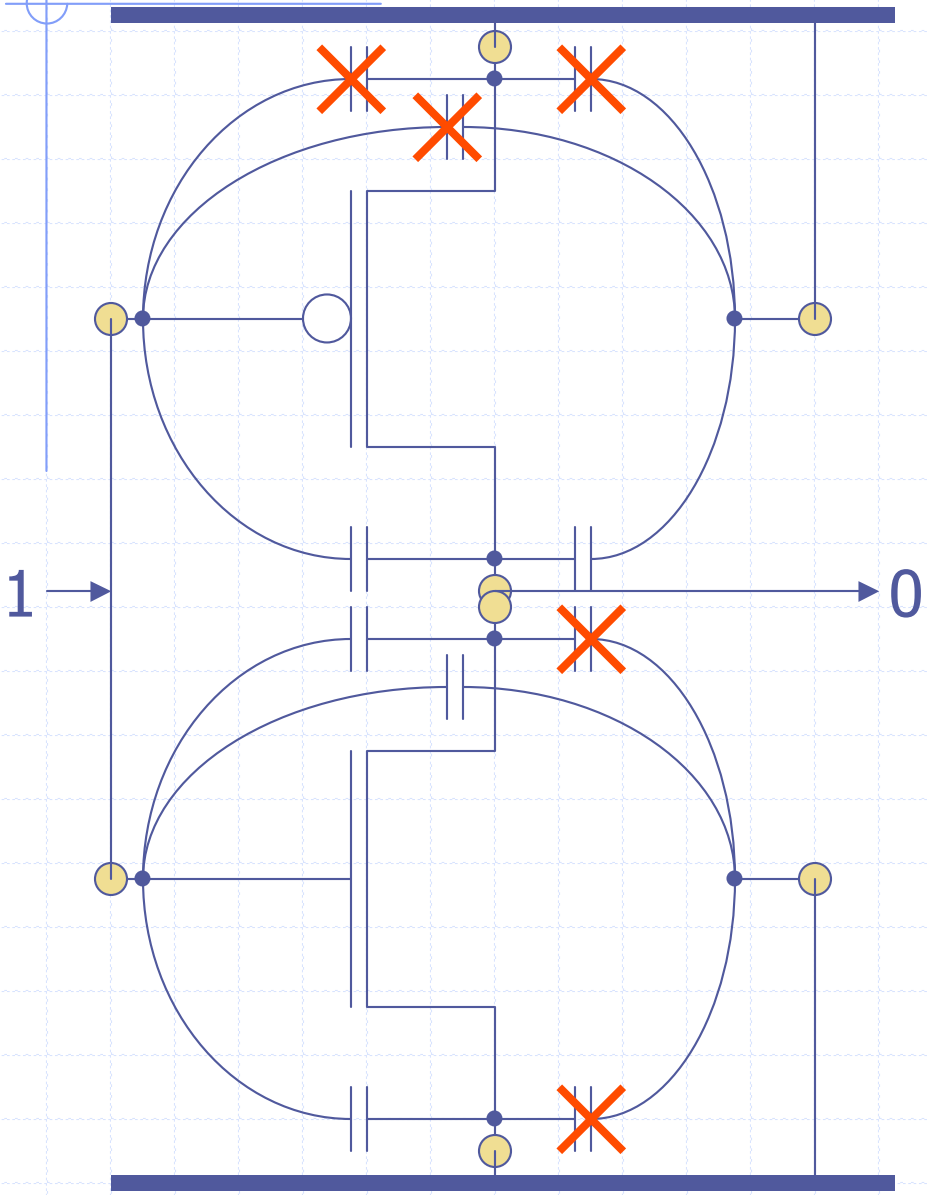

- With the input True, certain capacitors are discharged, while others are charged.
- The charged capacitors are the ones that can act as decoupling.
- Total decoupling capacitance depends on topology and on the state of the circuit.

Intentional Decoupling Capacitors

- Designers often add "intentional" decoupling capacitors.
- A possible design might use the C_{GX} components.
- Relatively easy to characterize, so will not discuss further in this work.

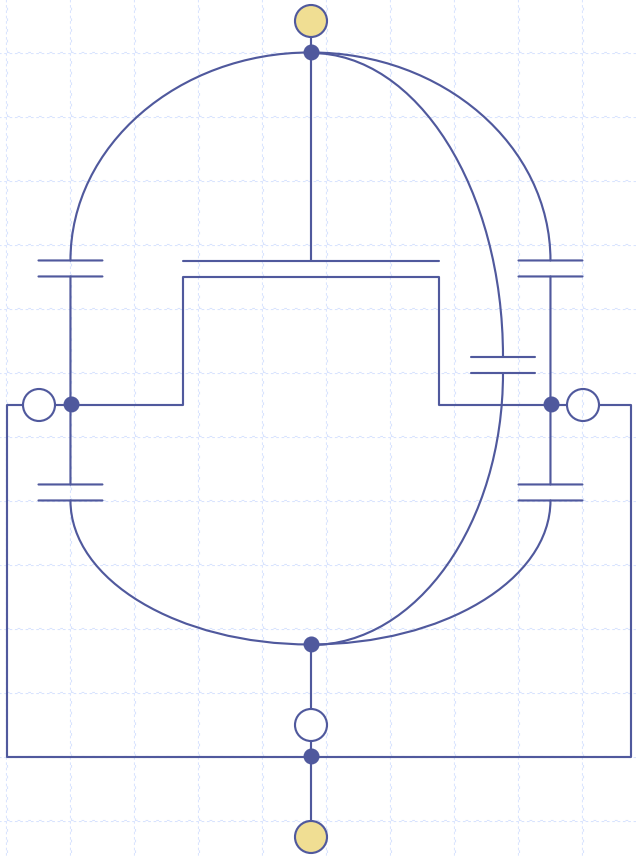

Decoupling Capacitance Estimation

- Heuristic method may rely on total device area (~ total capacitance) and de-rate it by some factor to account for the fact that certain components are charged / discharged.
  - $C_d = \alpha_N \Sigma A_N + \alpha_P \Sigma A_P$
  - A represents device area, α is the de-rating factor.
  - Summations over all devices.
- A more precise method can rely on simulation.

Simulation-Based Estimation

- Use AC analysis!

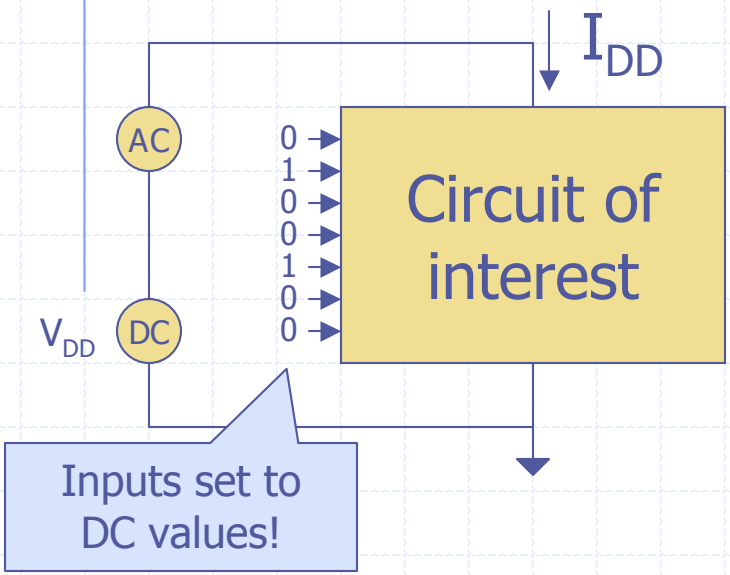

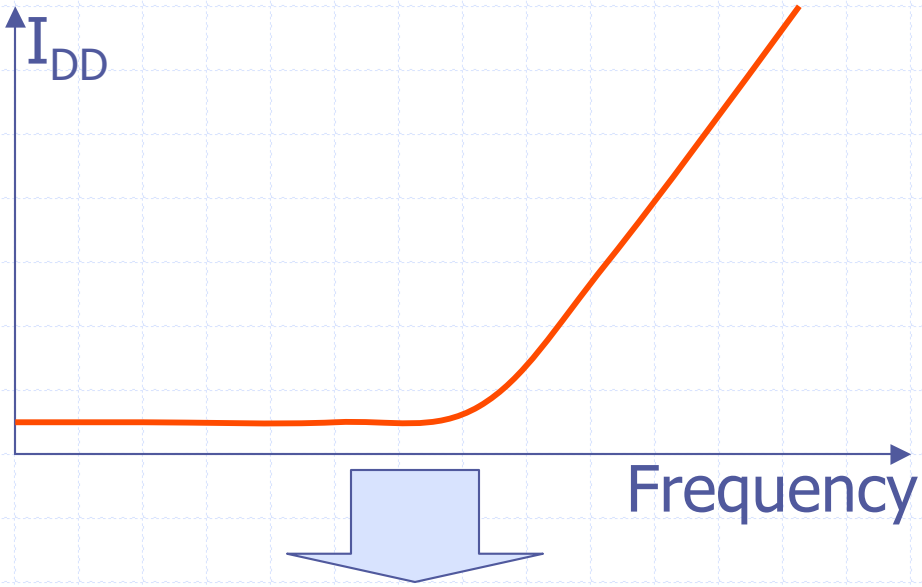

Equivalent RC circuit

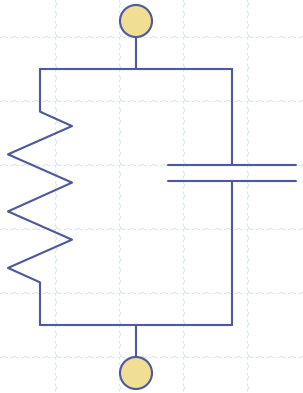

Simulation Example

- Same ISCAS C432 combinational circuit example.

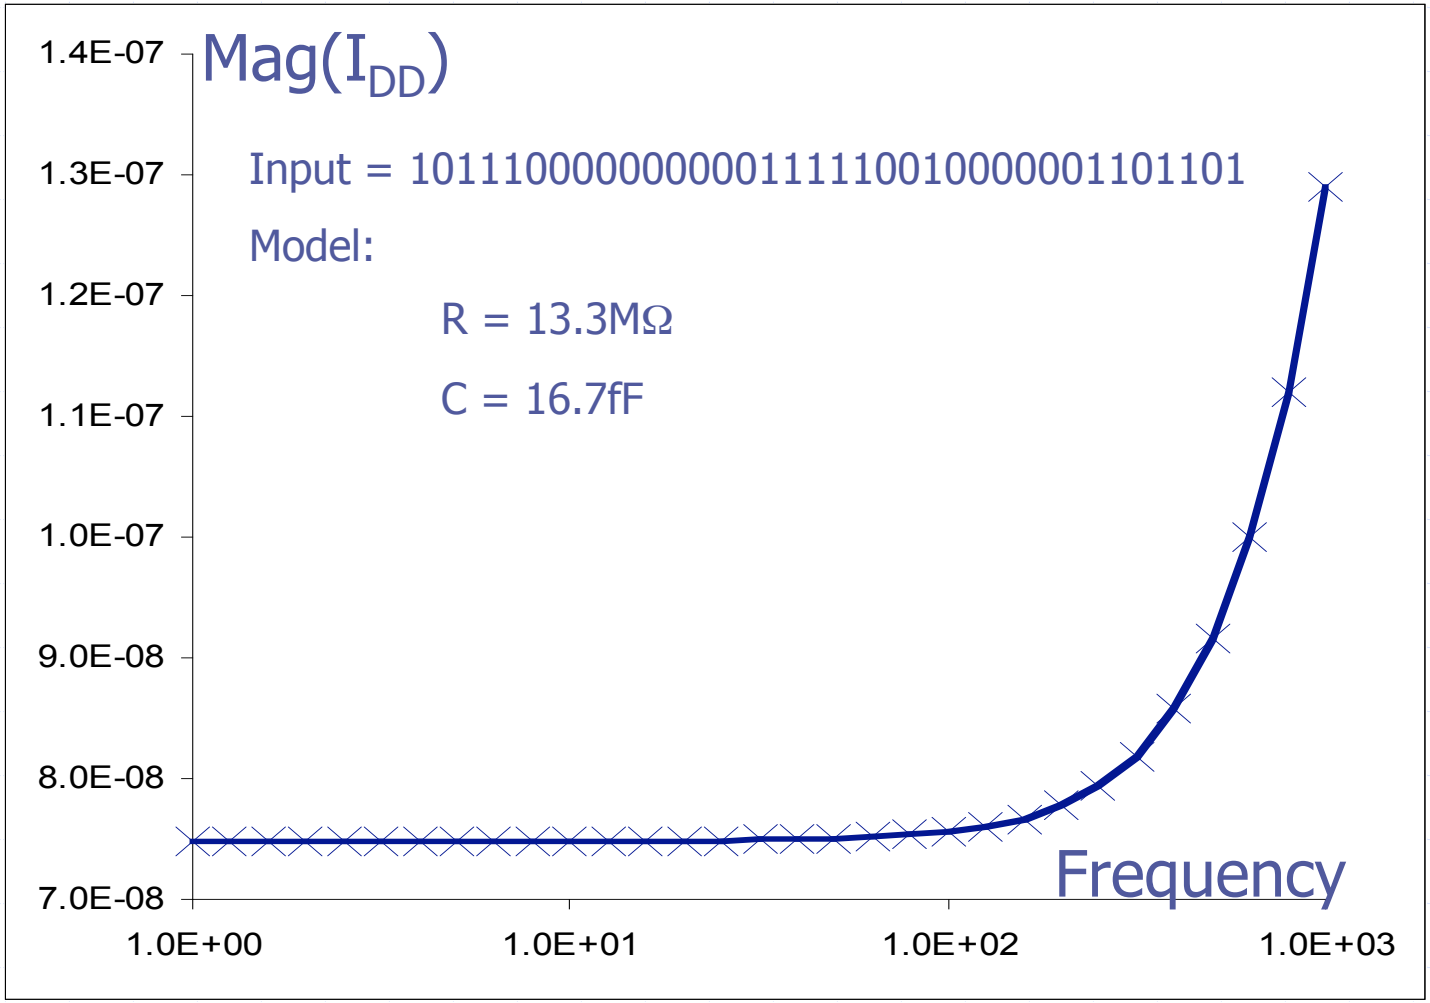

Pattern Dependence

- 64 different random input patterns, get R/C for each.

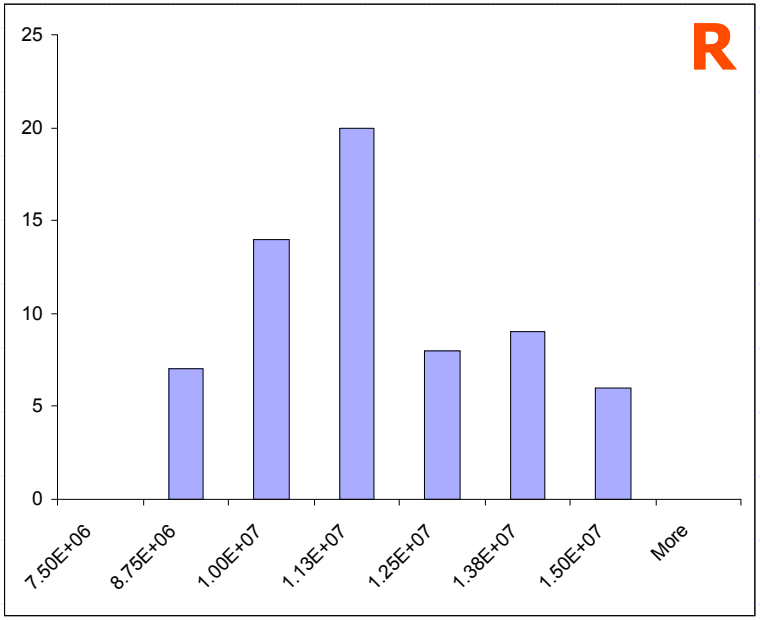

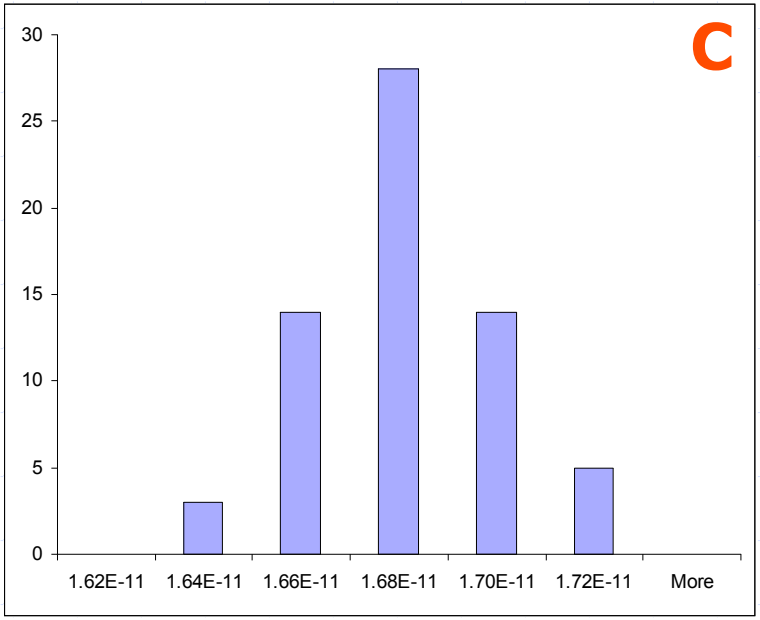

	R	C
min	7.87E+06	1.64E-11
max	1.49E+07	1.72E-11
μ	1.10E+07	1.67E-11
σ	1.87E+06	2.00E-13
σ/μ (%)	**17.0**	**1.2**

Very little pattern dependence!

Decoupling Capacitance Variability

- Variables of interest:
  - Technology
  - Circuit State (DC input pattern)
- Technology modeled via same collection of 58 sets of MOSIS 0.18μ parameters representing lot averages as used before.
- State modeled by a random sample of 64 unique input vectors.

Capacitance Variability Analysis

- Capacitance is approximately constant with respect to both process and pattern.
- Coefficient of variance (σ/μ) is less than 2%.

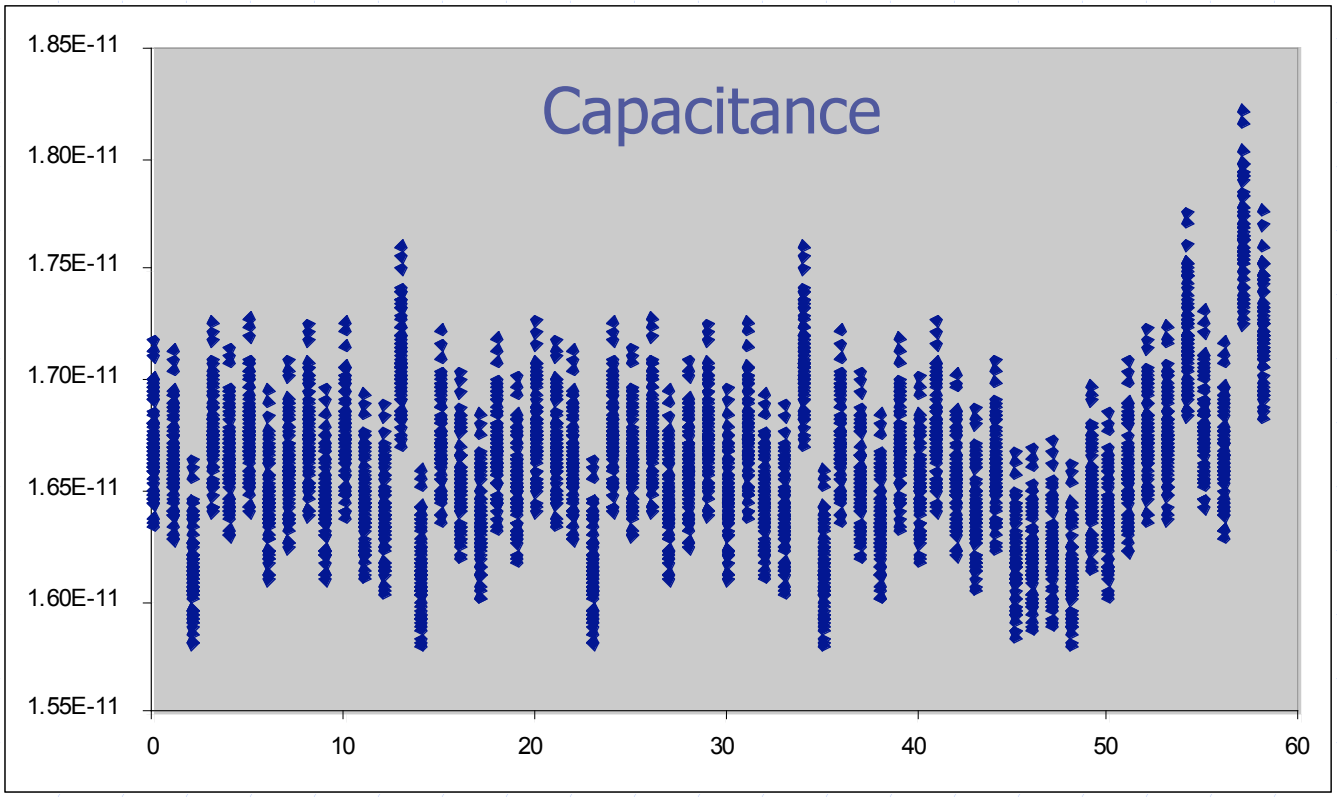

Resistance Variability Analysis

- Resistance varies substantially with respect to both process and pattern.
- Note that Resistance ≡ Leakage!

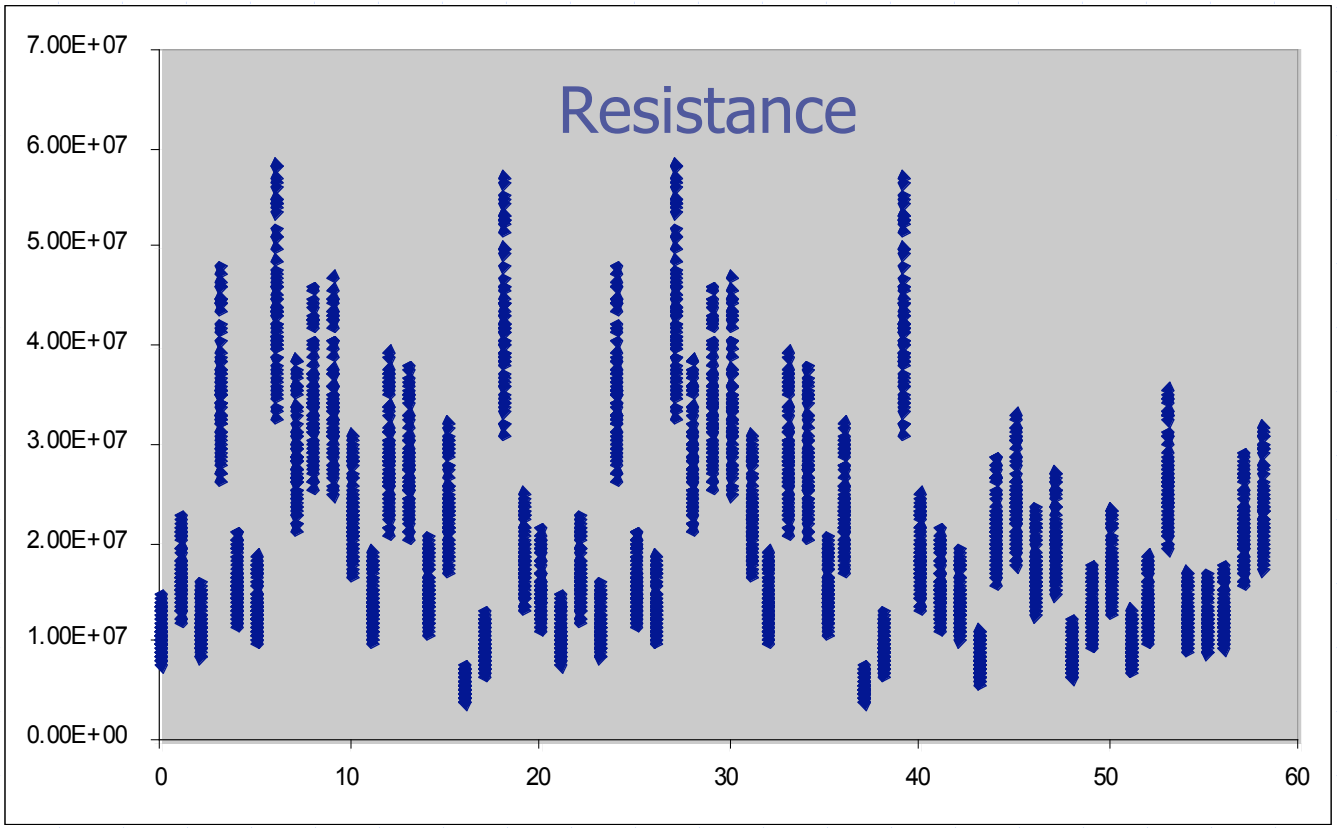

Resistance Variability Analysis

- Compare means vs. process and vs. pattern.

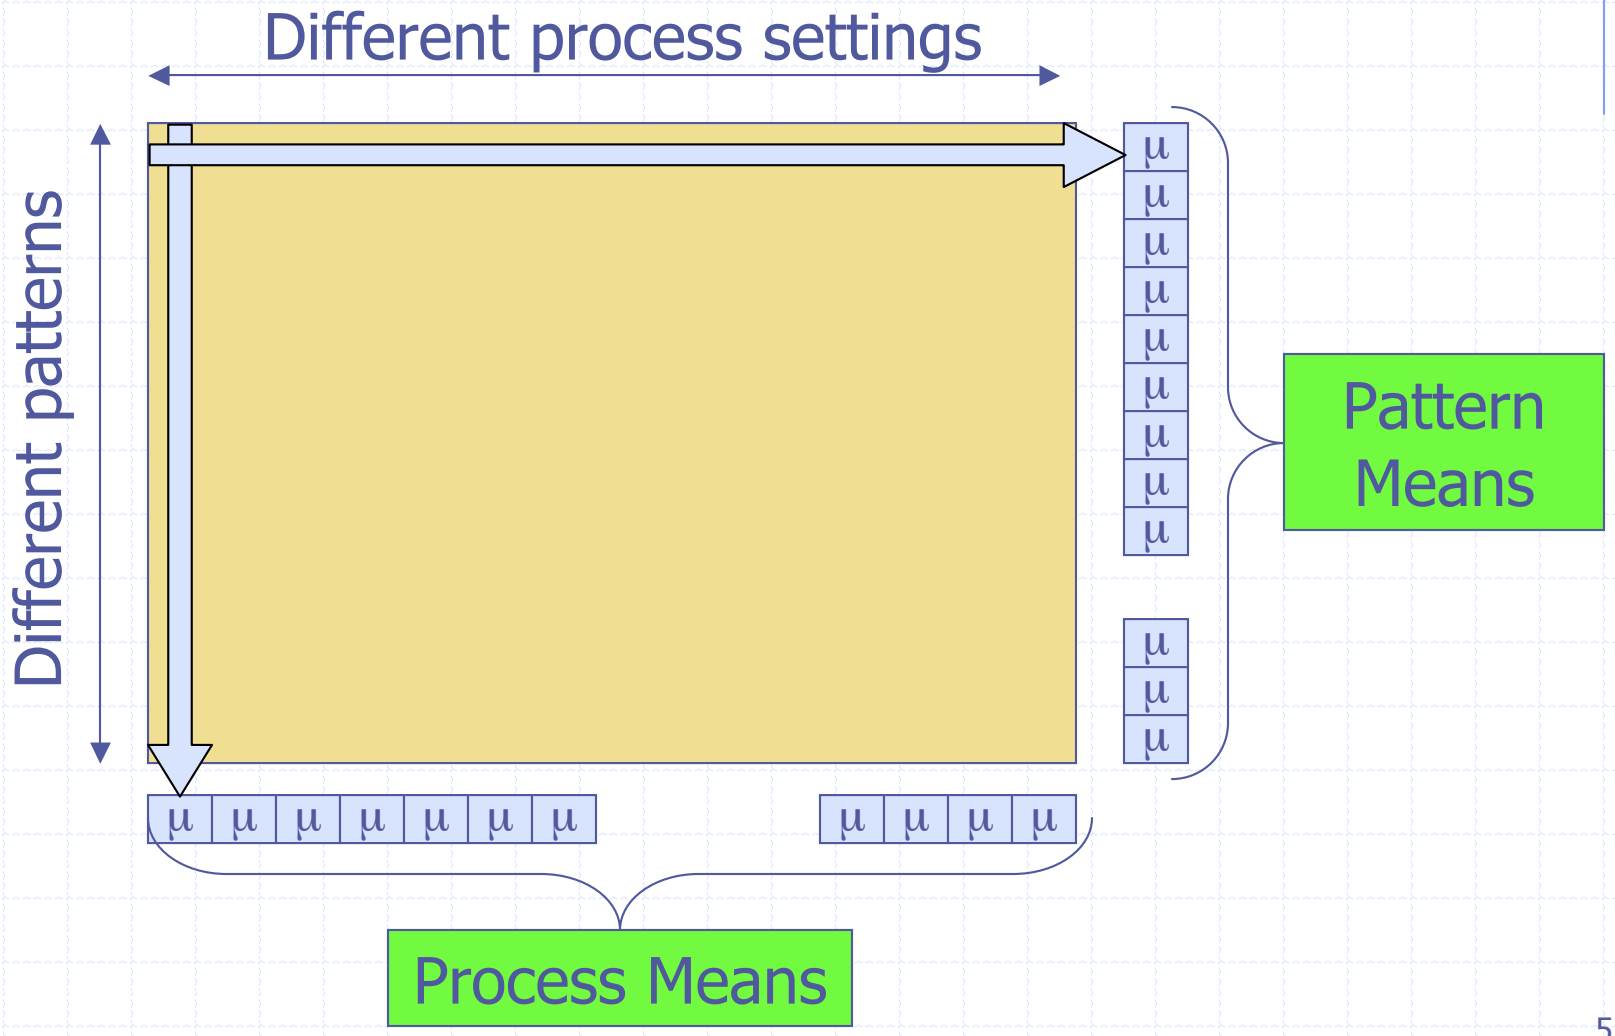

Process vs. Pattern Means

Pattern

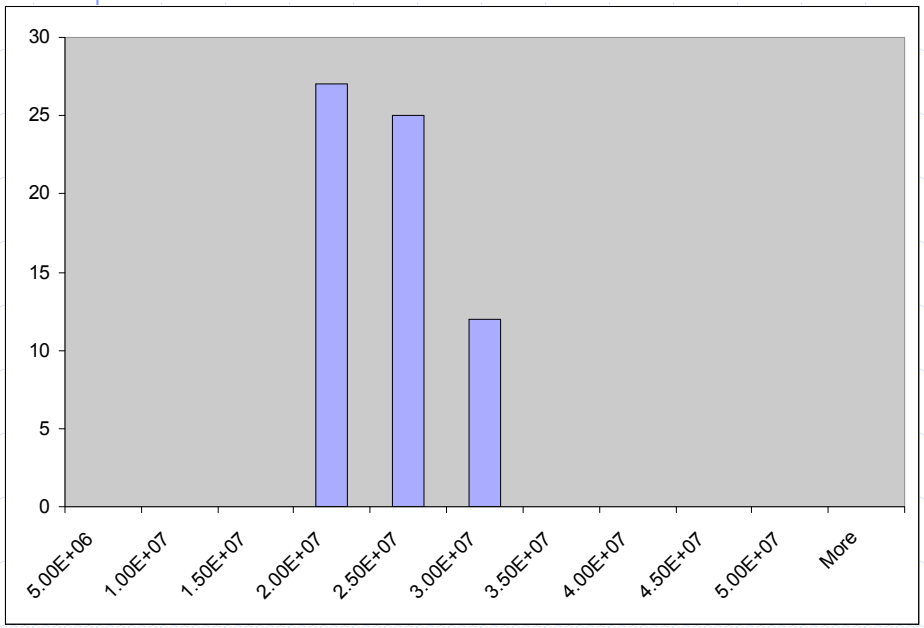

Process

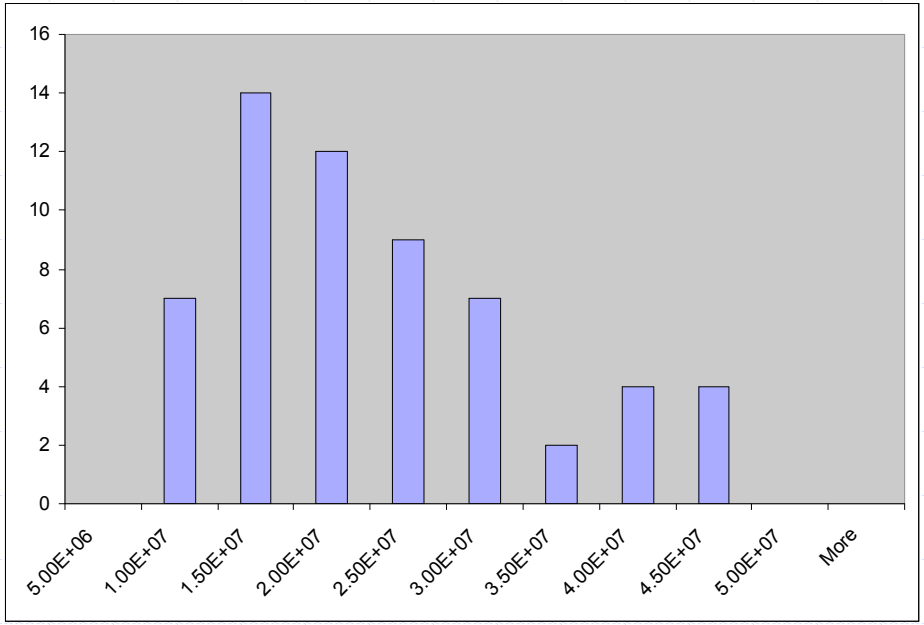

5x more variability due to process

Conclusions

- Decoupling capacitance can be estimated using a simple AC analysis.
- The capacitive part of the decoupling capacitance of a circuit is ~ constant.
- The resistive part of the decoupling capacitance of a circuit depends strongly on technology and weakly on circuit state.
  - For certain types of analyses, however, the resistance may not be needed.
  - BUT… Resistance variations are useful to assess leakage variations.

Outline

- Introduction
- Variability & Uncertainty
- Power Delivery Components
  - Circuits (power dissipation)
  - Decoupling Capacitance
  - On-Chip Power Grid
  - Package
- Tool Requirements
- Conclusions

On-Chip Power Grid

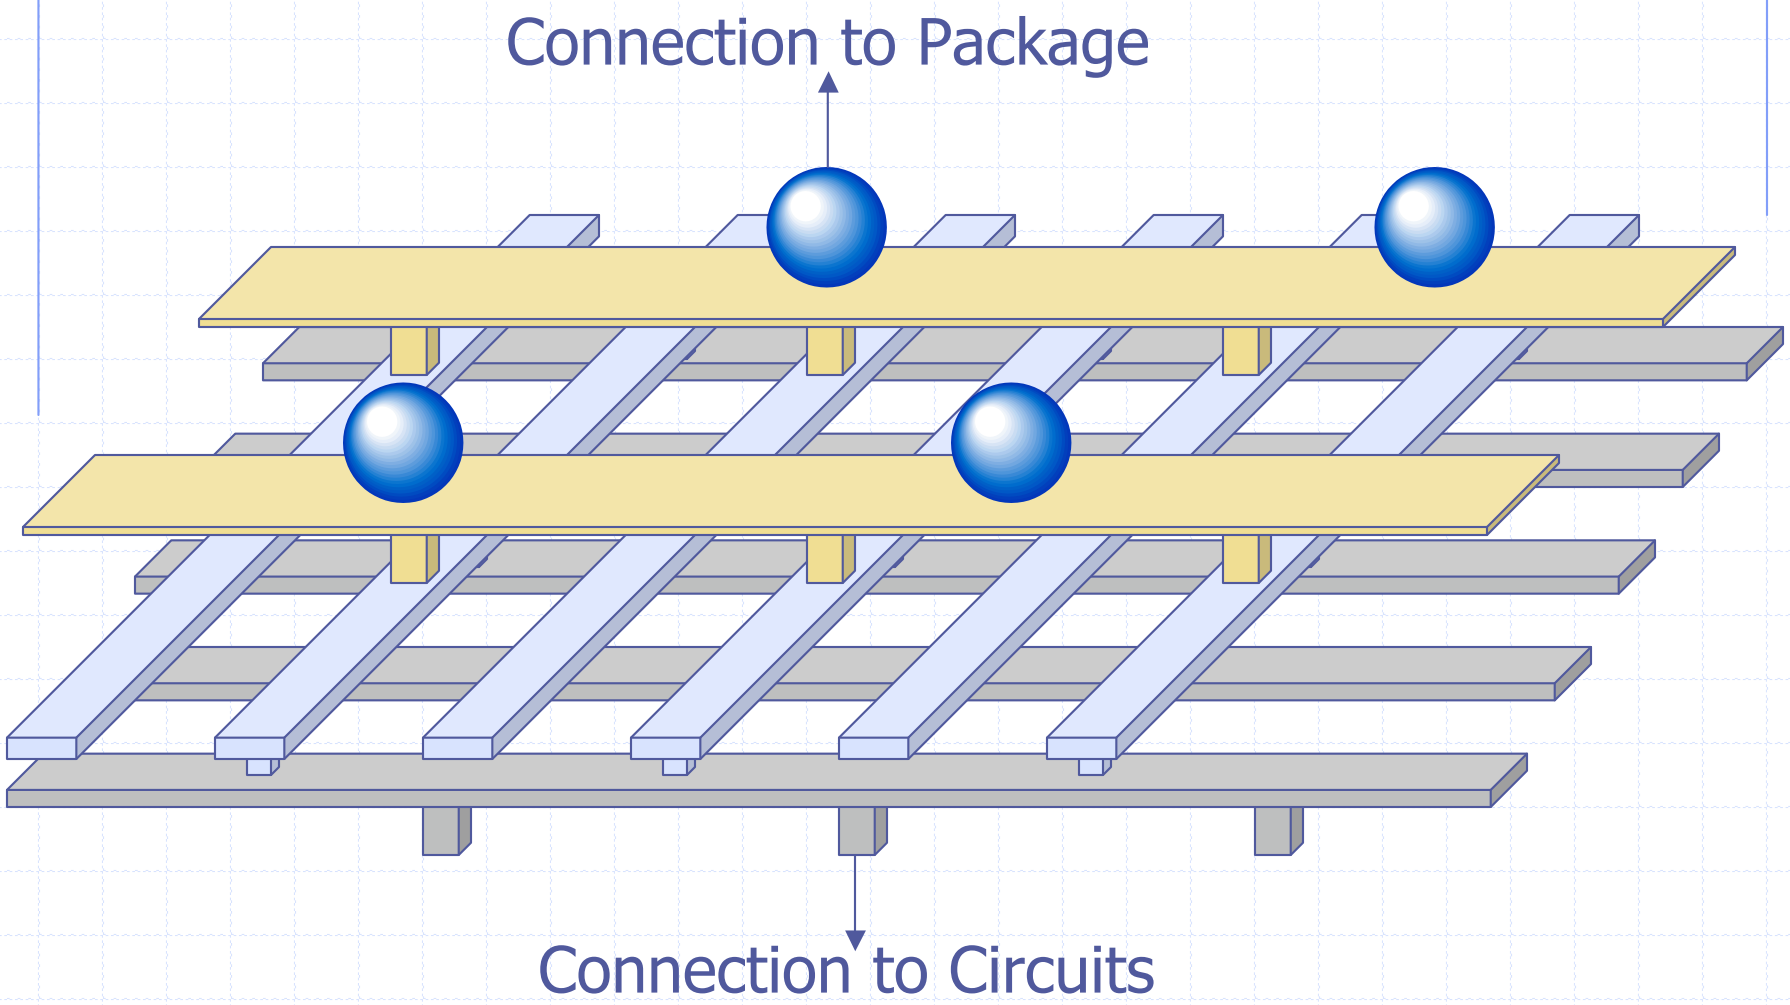

On-Chip Power Grid Variability

- Determine contribution to IR drop variability of various grid system components:
  - Grid wires.
  - Package resistance.
  - Spatial power consumption (design).

Methodology:

- Perform a Design Experiment, measure total voltage drop (differential) and create a linear model with respect to design parameter.

Design Experiment

- 3 x 3 C4 region.
  - ~1500 x 1500 μ.
- 6 levels of metal
  - Two at 1x, two at 2x, and two at 4x thickness.
  - Both VDD and GND wires included.
- Nominal drop of 10% of V_{DD} at the center.

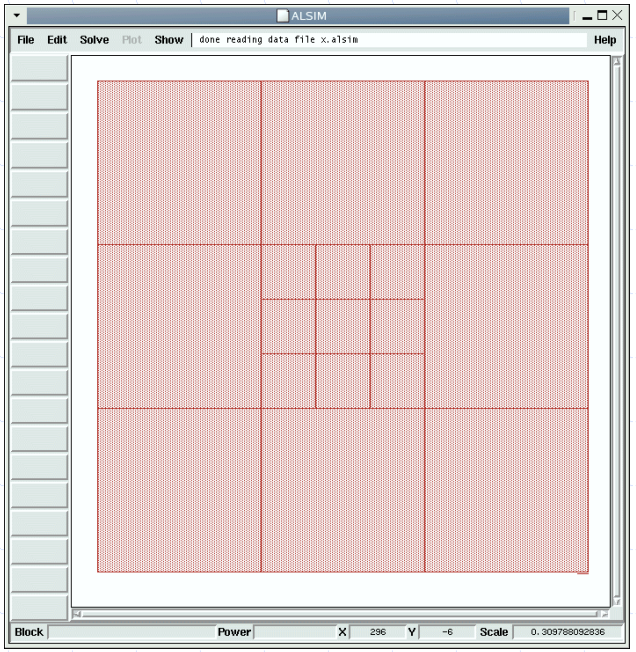

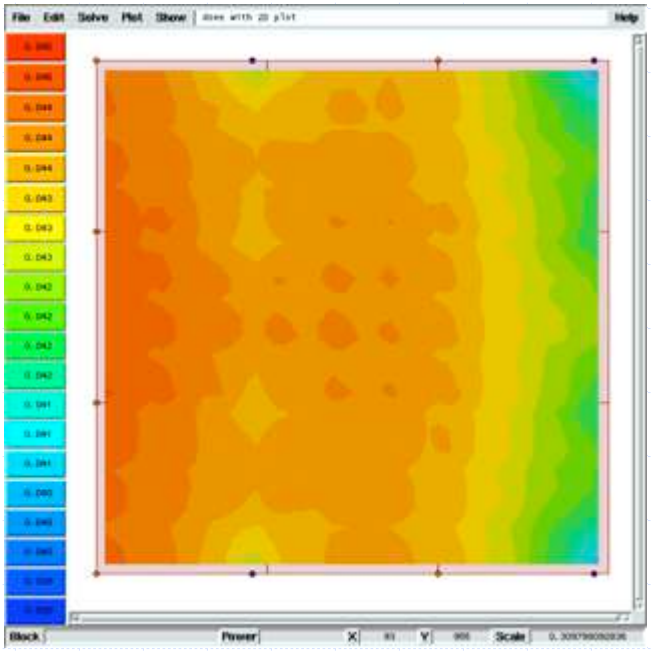

Power Grid Wires

- M1, M2 15% density $\rho=0.08\Omega/\square$
- M3, M4 20% density $\rho=0.04\Omega/\square$
- M5, M6 25% density $\rho=0.02\Omega/\square$

- Used a 10% tolerance in grid resistivity.
  - Since the goal is to build a model, exact values are not important. Relative importance of the various effects is.
- Package resistance taken as 0.25Ω per pin with a 10% tolerance (remember for later…).

Power Grid Loading

- Three variables in experiment:

1. Loading (DC) at test point.
2. Loading in local loading blocks.
3. Loading in global loading blocks.

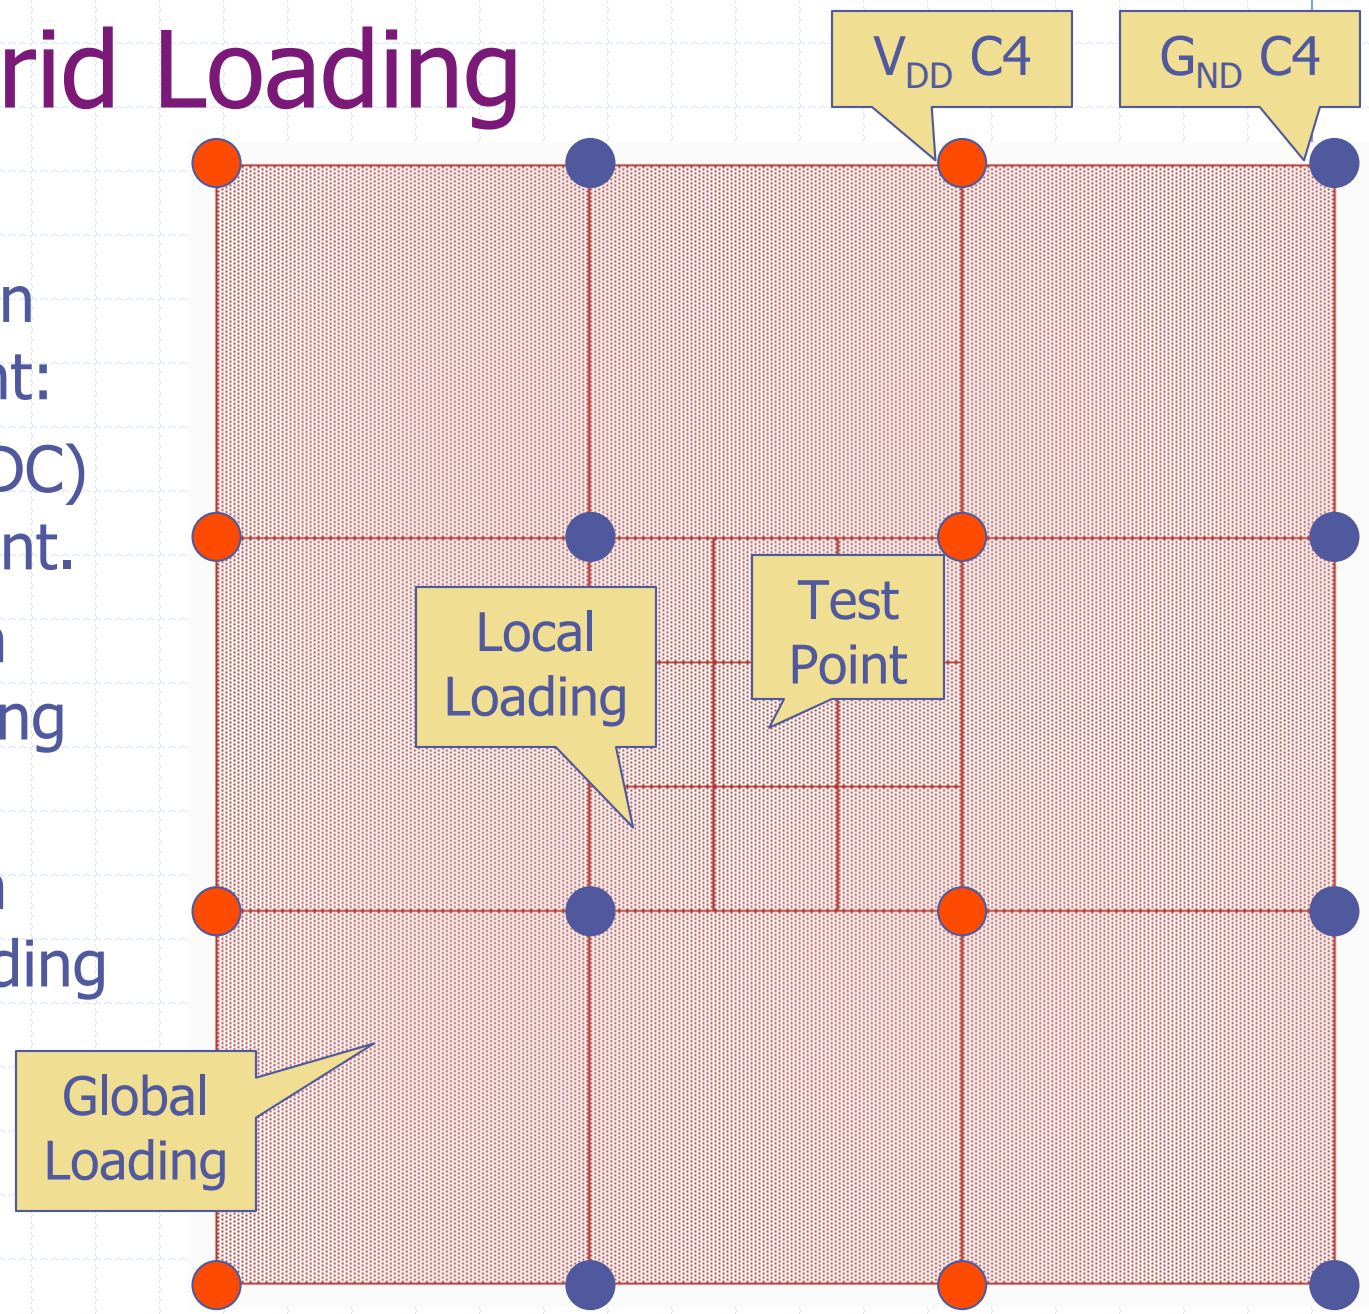

Example Output: V_{DD} Distribution

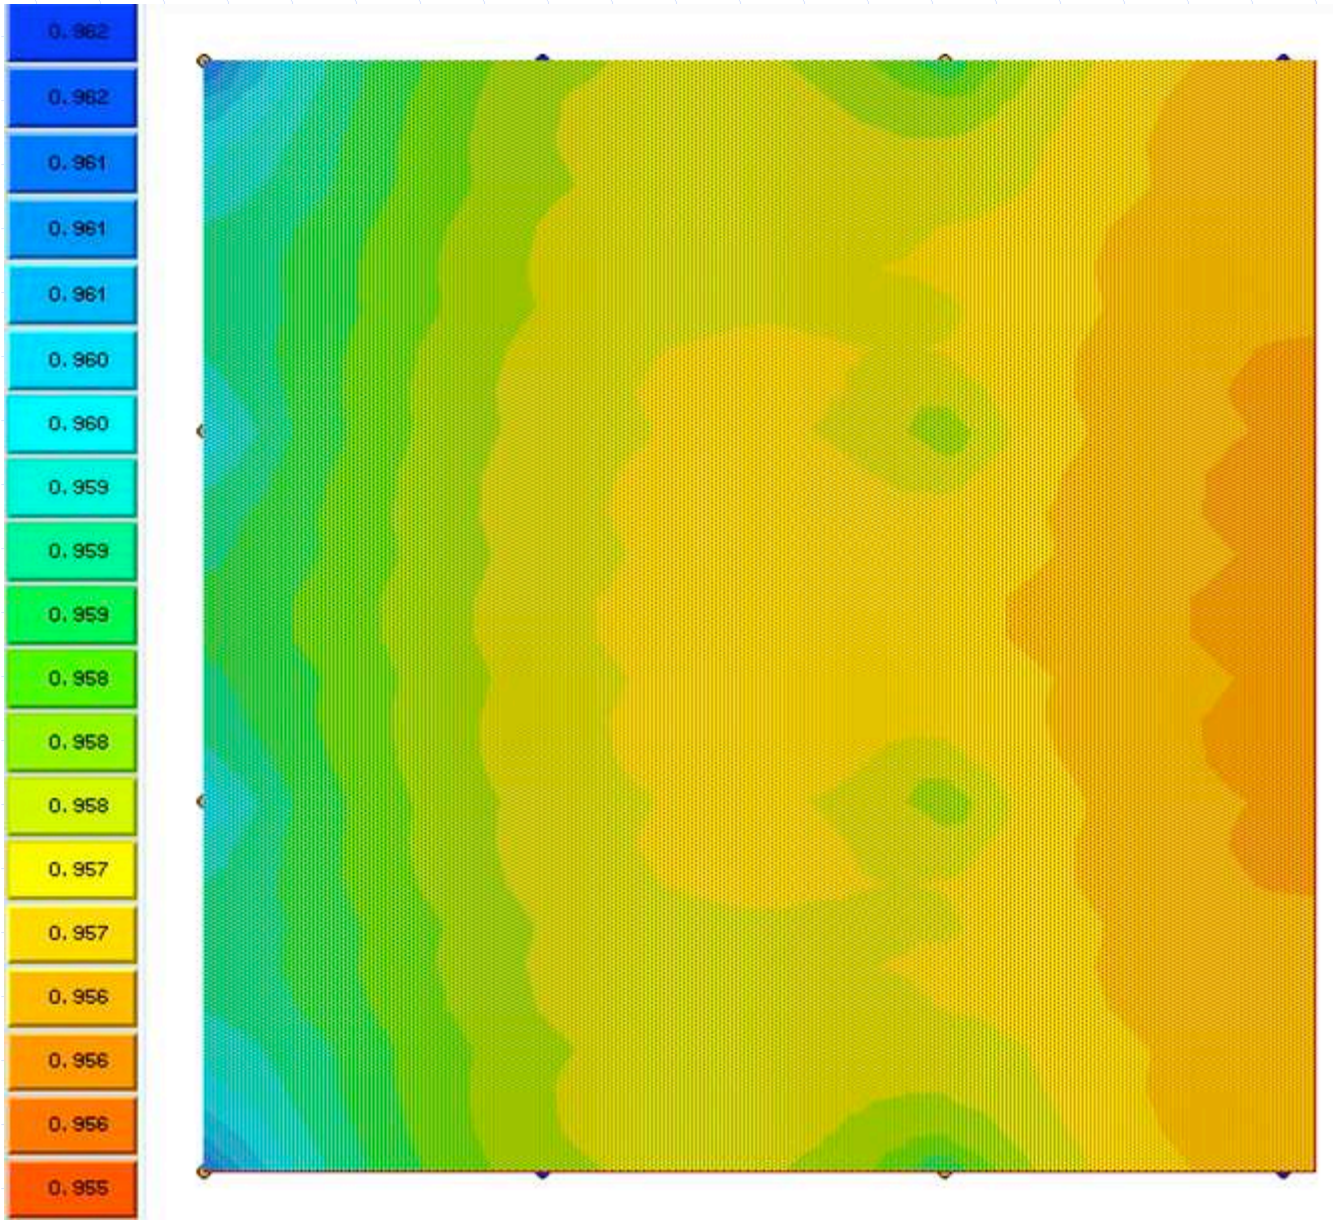

Experiment Structure

- Used Latin-Hypercube sampling to generate a sample of 200 uniformly distributed parameter settings for:
  - Metal sheet resistivities.
  - Local and Global loading parameters.
  - Package per-pin resistance.
- Larger sample sizes produced essentially the same results!

Results: Model Fit

Modeled V_{DD} at center as a linear function of the various design variables.

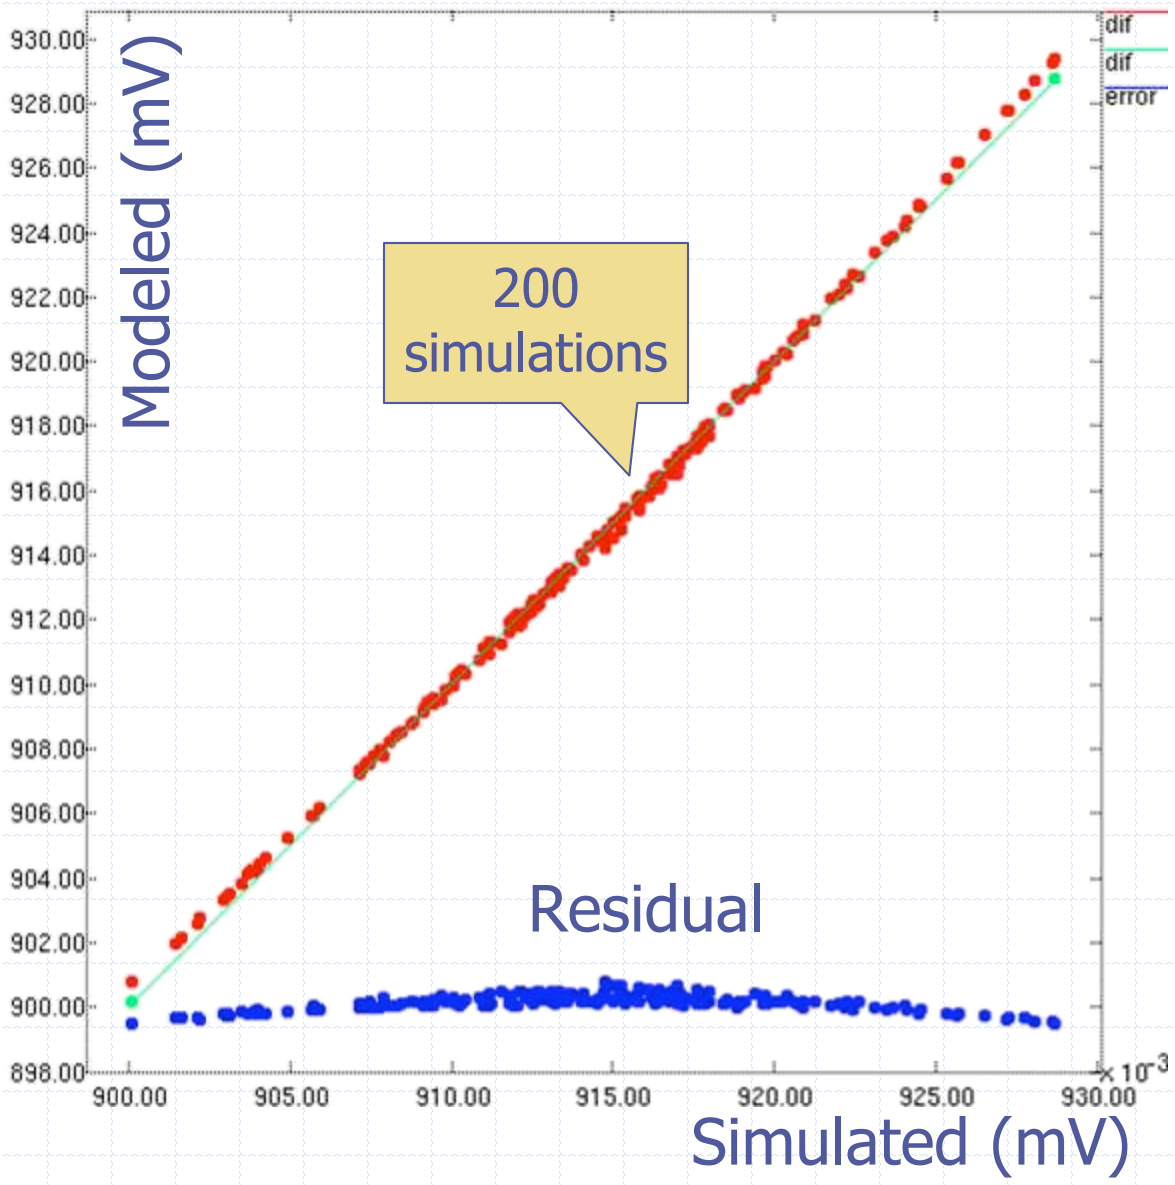

Results: Model Coefficients

Variable	Coefficient	Nominal	Normalized
M1	0.000994	0.08	0.00008
M2	0.008102	0.08	0.00065
M3	0.006799	0.04	0.00027
M4	0.003414	0.04	0.00014
M5	0.054165	0.02	0.00108
M6	0.076293	0.02	0.00153
Package	0.323526	0.25(16)	0.00506
B-center	0.593441	0.002	0.00119
B-local	0.504067	0.002(8)	0.00101
B-global	0.058480	0.002(72)	0.00012

Relative Impact

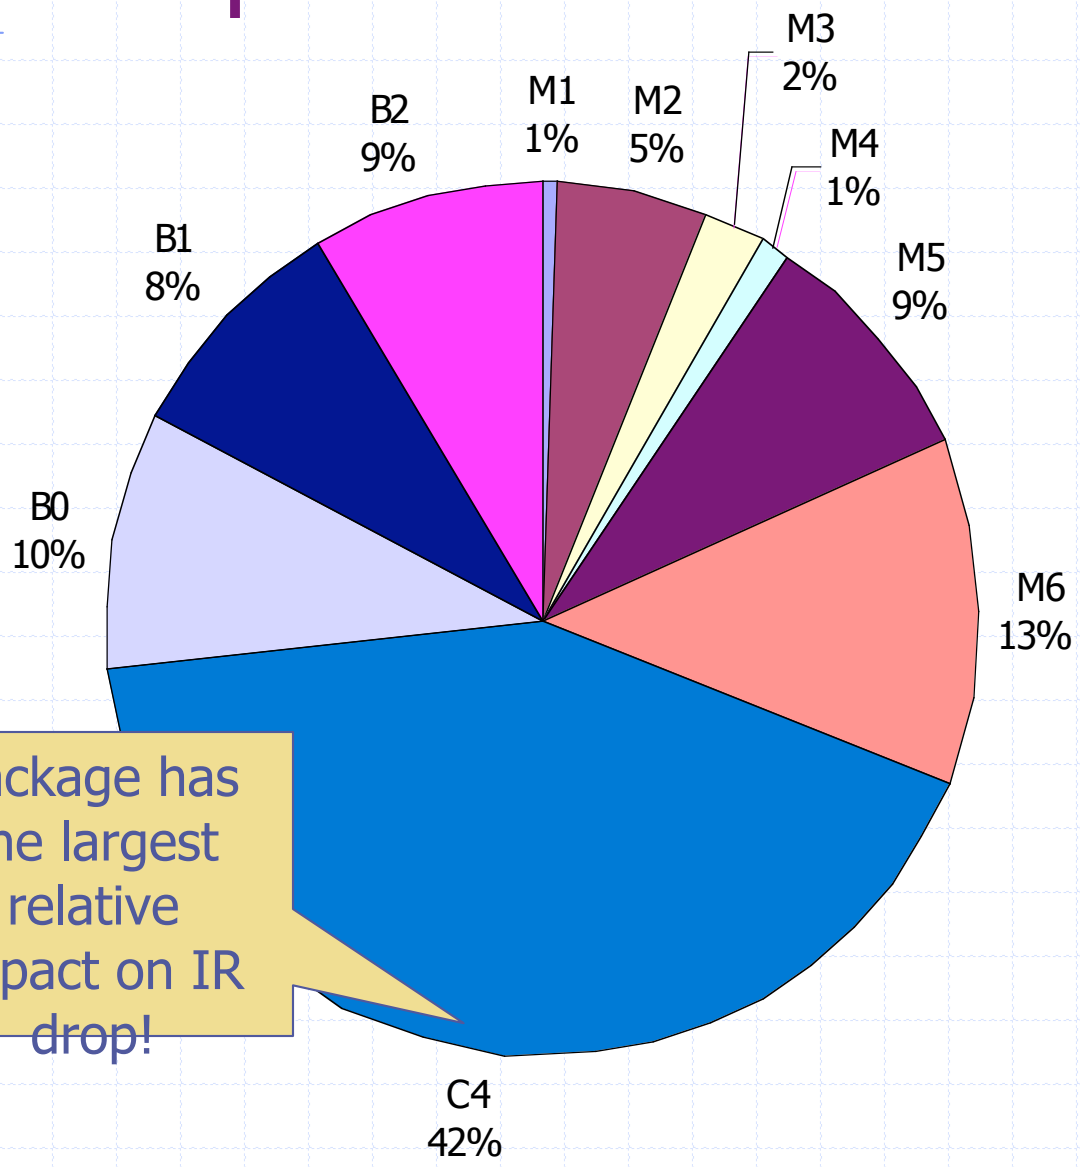

Conclusion

- Power grid variability (DC) can be readily assessed by performing a designed experiment.
  - A similar analysis can be done for AC.
- Lower levels of metal are less variation sensitive than higher levels.
- Package parasitics play an important part in the overall variability.

Outline

- Introduction
- Variability & Uncertainty
- Power Delivery Components
  - Circuits (power dissipation)
  - Decoupling Capacitance
  - On-Chip Power Grid
  - Package
- Tool Requirements
- Conclusions

Package Variability

- Packages are composed of wiring planes, sometimes with embedded discrete decoupling capacitors.
- Dimensions are such that R and L are important, but C is not.
- To first order, manufacturing variability in R is ~ resistivity. Variability in L is much smaller.
- BUT… packages are rarely symmetric so the systematic variability from one pin to the other becomes dominant!

Package Variability Estimation

- Used an in-house tool to extract equivalent resistance and inductance for each C4 of a typical mid-range ASIC package.
- Tool is based on detailed layout extraction and uses an L^{-1} formulation to perform accurate full-package inductance extraction.
- Processing time for full package: ~3 hours.
  - This is a **very hard problem!**

Example: Package Statistics

- Package Size : 23mm X 23mm
- Chip Size : 10mm X 10mm
- Layers : 17
- Top Pins : 129 VDD + 261 GND
- Bottom Pins : 36 VDD + 80 GND
- # of shapes : ~ 80,000
- R extraction time : 0.5 min
- L extraction time : 156 min

V_{DD} Pin Systematic Differences

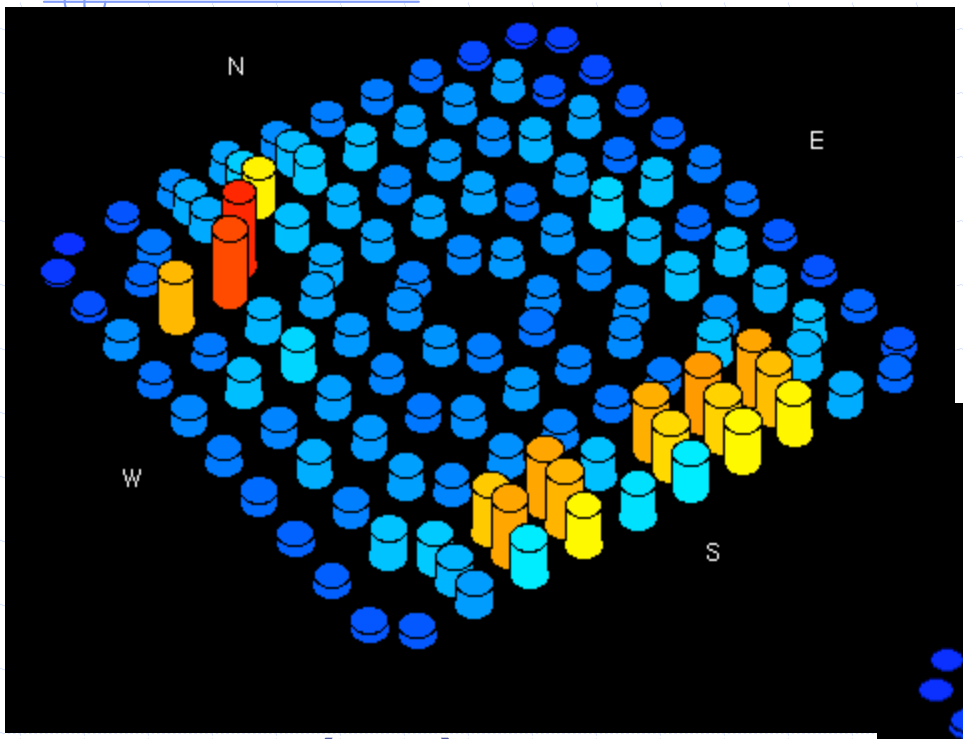

resistance (mΩ)

inductance (pH)

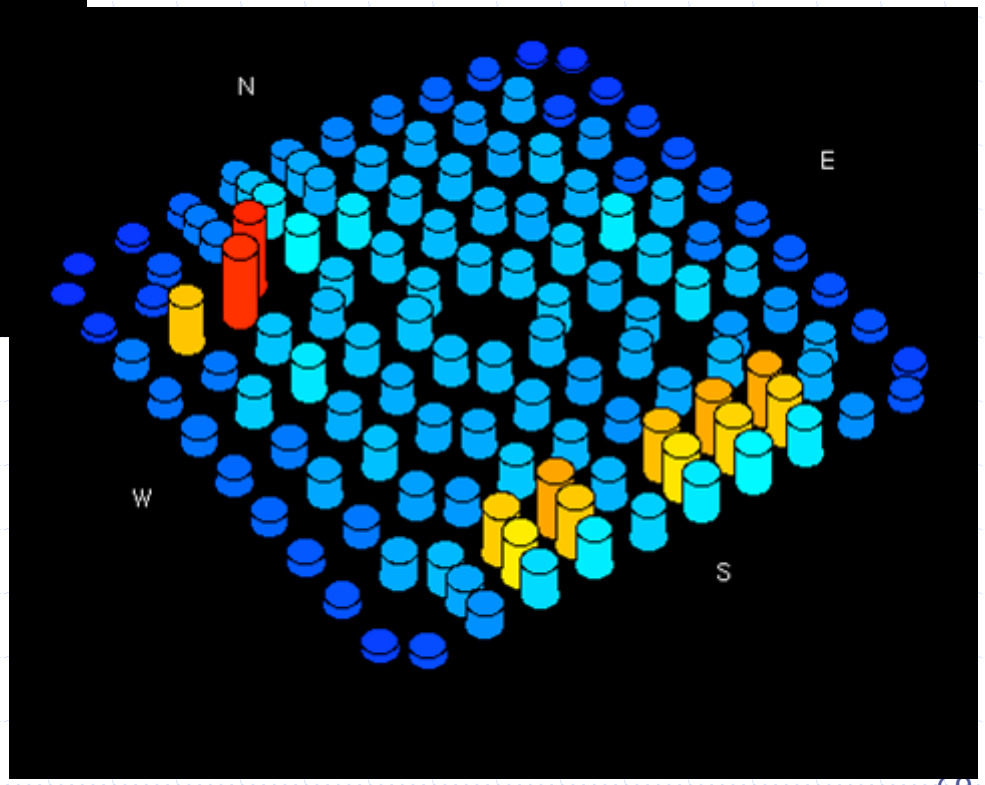

Systematic Resistance Variations

Statistics (mΩ):

- μ=263.4
- σ=63.6 (24%)
- Min=157.5
- Max=519.6

- A 10% tolerance on resistivity is insignificant compared to the systematic variations!

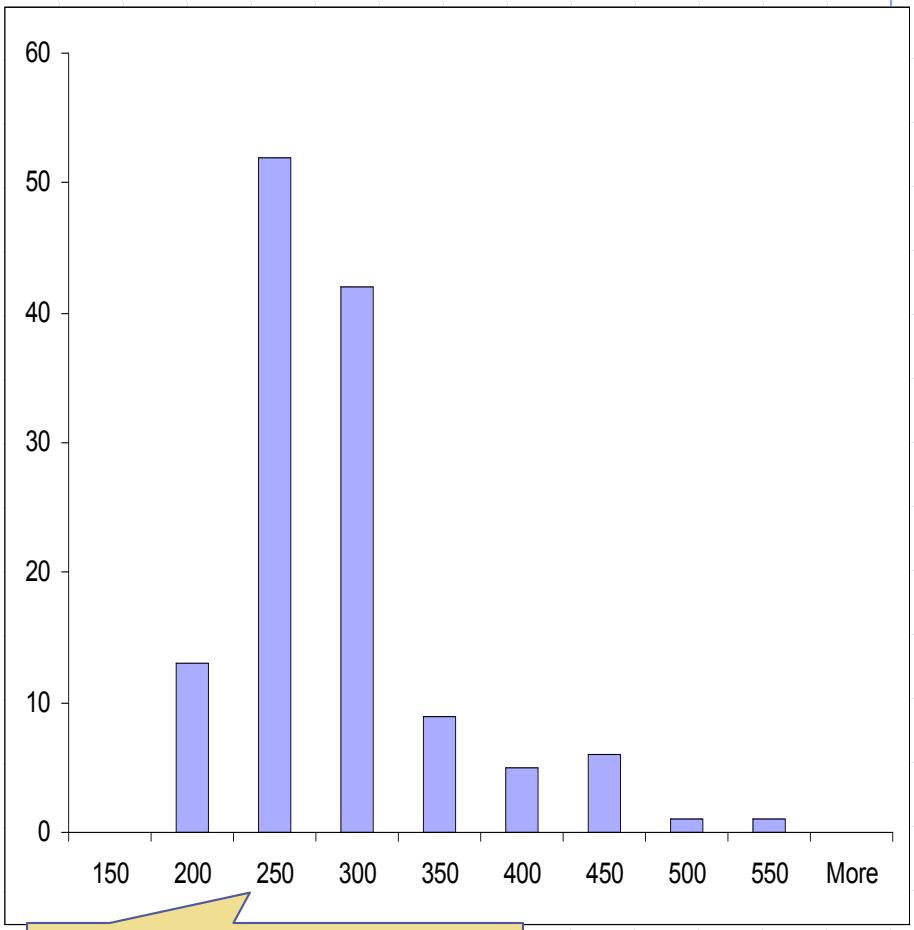

Remember that we used 0.25 Ω ± 10% in the Power Grid Experiment!

Systematic Inductance Variations

Statistics (nH):

- μ=4.24
- σ=0.57 (13.5%)
- Min=3.11
- Max=6.61

- Inductance has a smaller coefficient of variation than resistance!

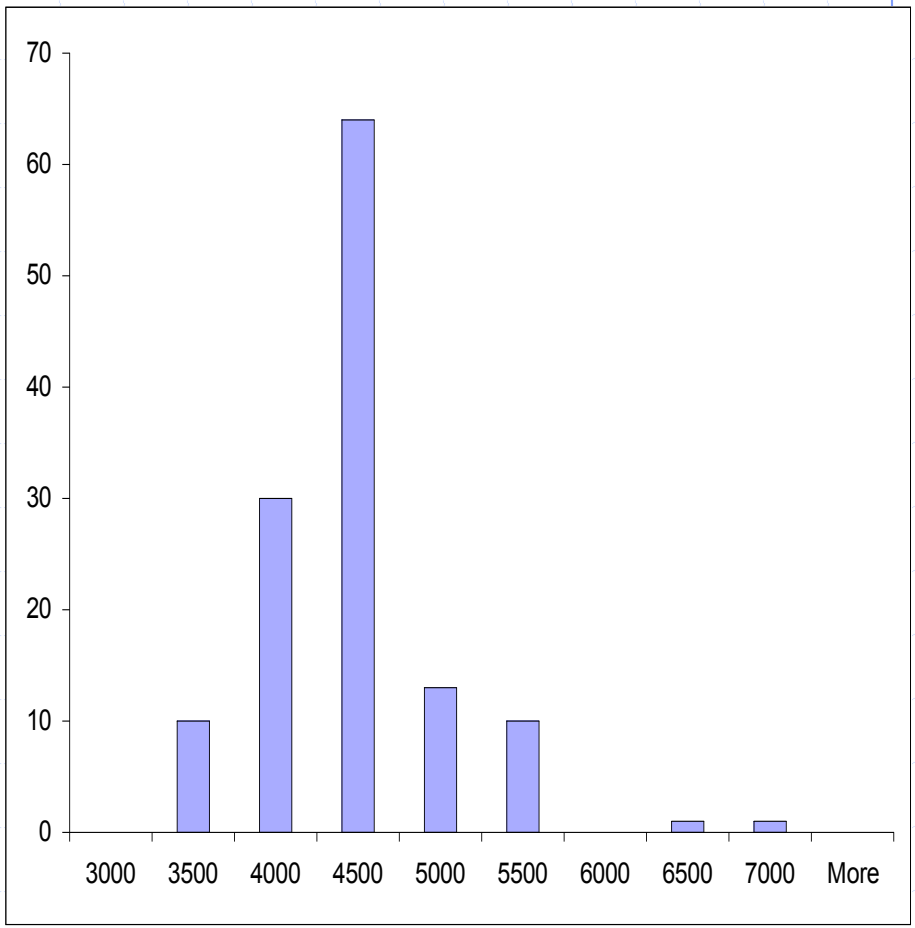

GND Pin Systematics: Resistance

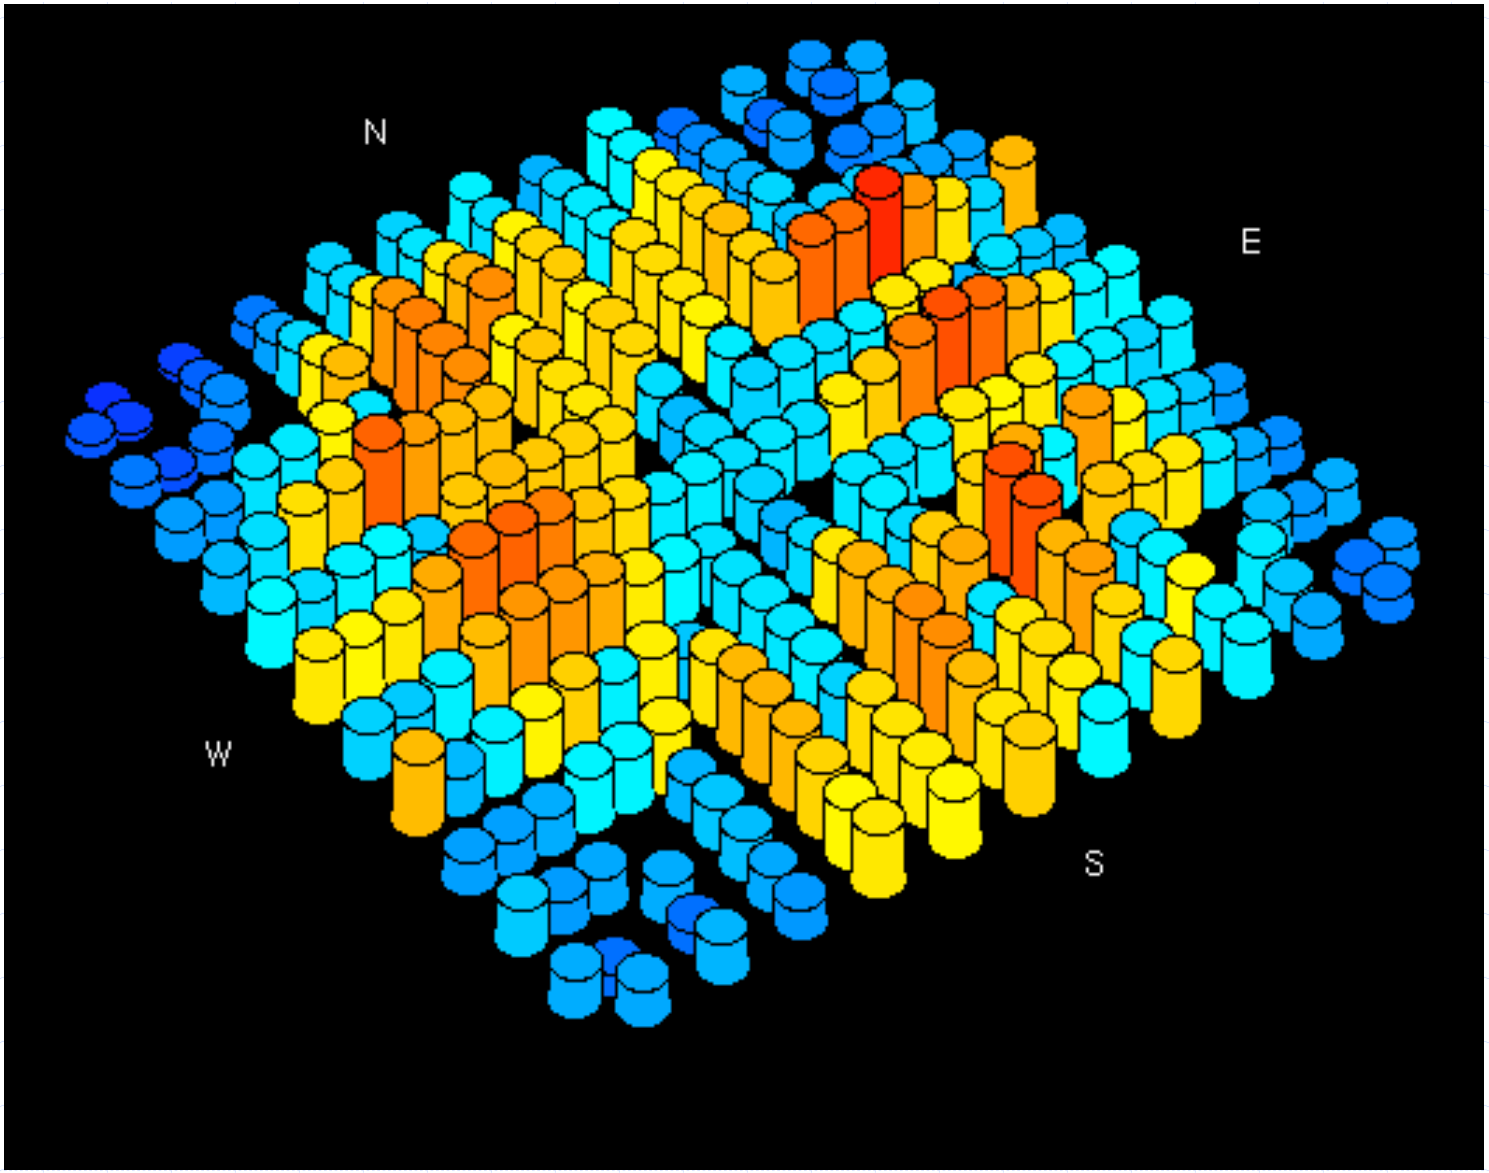

GND Pin Systematics: Inductance

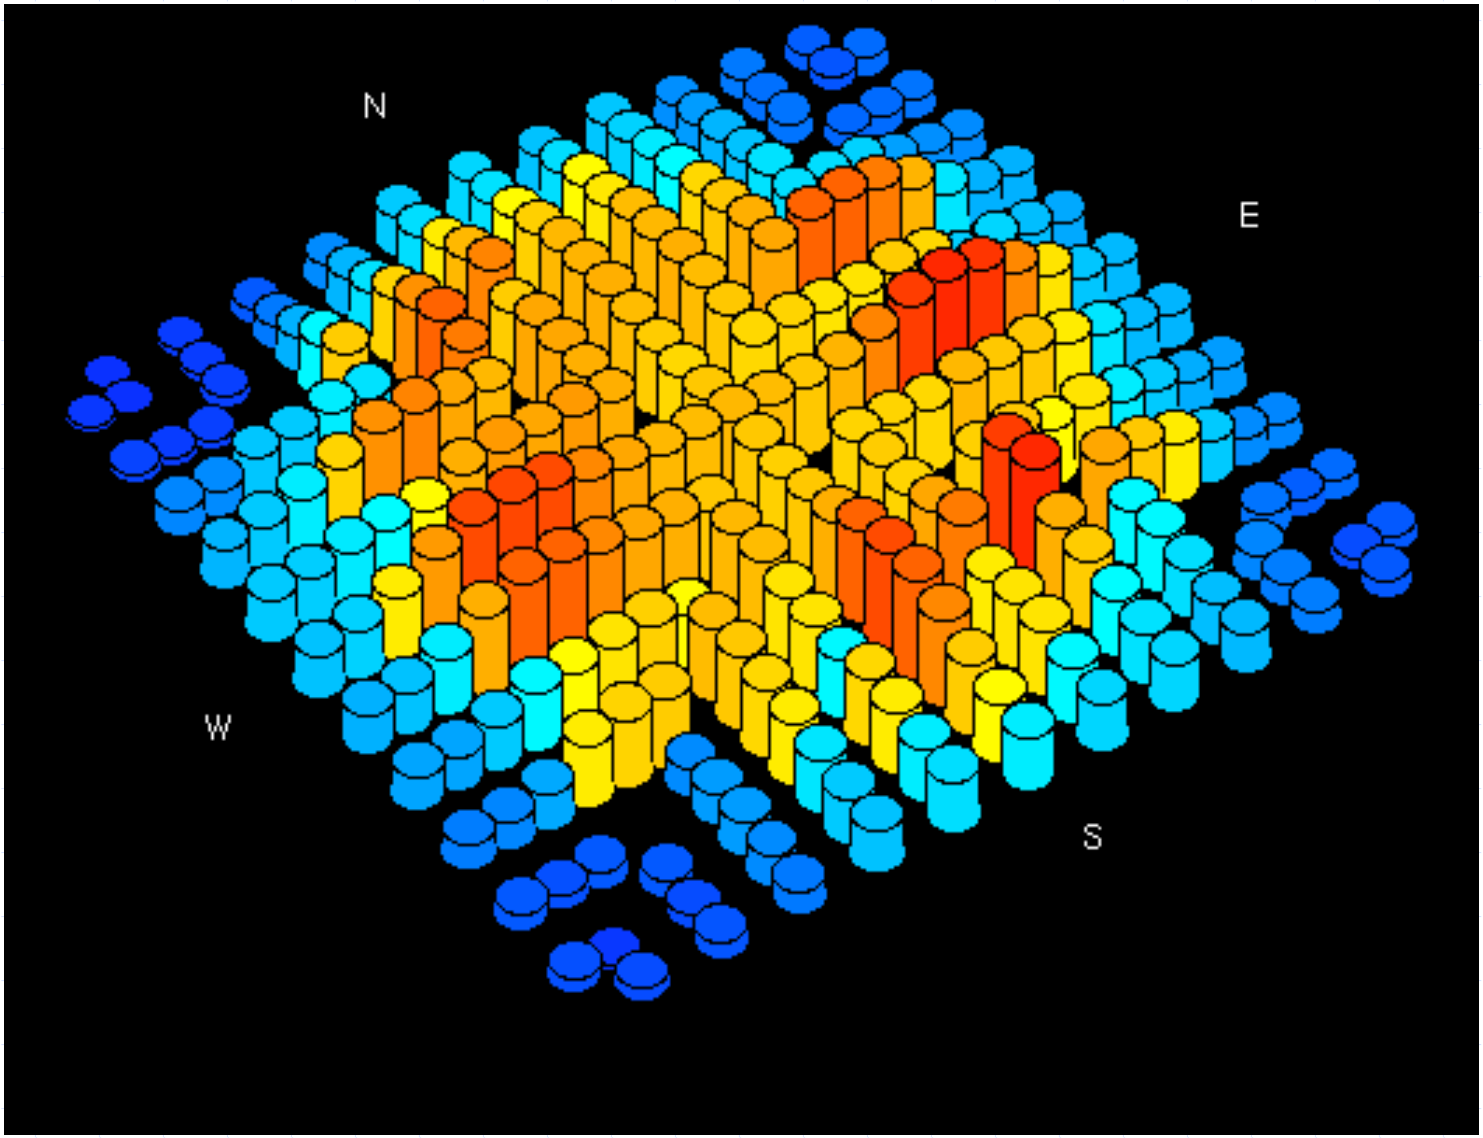

Conclusion

- Package parasitics (for high performance packages) have systematic variability that is large compared to expected manufacturing tolerances.
- Not being able to assess such variability makes it an uncertainty and may cause excessive over-design!
- The lack of automated full-package analysis tools is a huge problem in this area.

Outline

- Introduction
- Variability & Uncertainty
- Power Delivery Components
  - Circuits (power dissipation)
  - Decoupling Capacitance
  - On-Chip Power Grid
  - Package
- Tool Requirements
- Conclusions

Power Grid Design Trends

- Increasing number of levels of metal.
  - More degrees of freedom for tradeoff between interconnect and power.
  - More effort in grid design.
- More design restrictions on wires (Cu, CMP).
  - Example: maximum width, metal density, oxide density within metal area, etc…
- Package design choices critical:
  - More package power pins (fixed I_{MAX} per pin).
  - Area array (distributed) versus wirebond (edge).
  - May need to correct for systematic errors on chip.

Power Grid Design Issues

- Power grid design:
  - Done before detailed implementation starts.
  - Spatial power requirements approximated.
  - Impacts implementation of all Physical Design components.
    - Placement of high-power devices (I/O, clock, arrays).
    - Placement and allocation of decoupling caps.
  - "Interface" between power distributions costly.
- Result:
  - Rampant over-design.
  - 15 to 20% of wiring resources needed.

IBM Power Grid Planner

- To help explore options early in design cycle.
- Tool needs to be very fast (interactive).

Typical questions:

- Can a grid with X% metal density handle P watts per square mm?
- How much decoupling cap does an I/O buffer need? How close does it need to be?
- How much reduction in L di/dt do I gain by introducing deliberate skew?

IBM Power Grid Planner

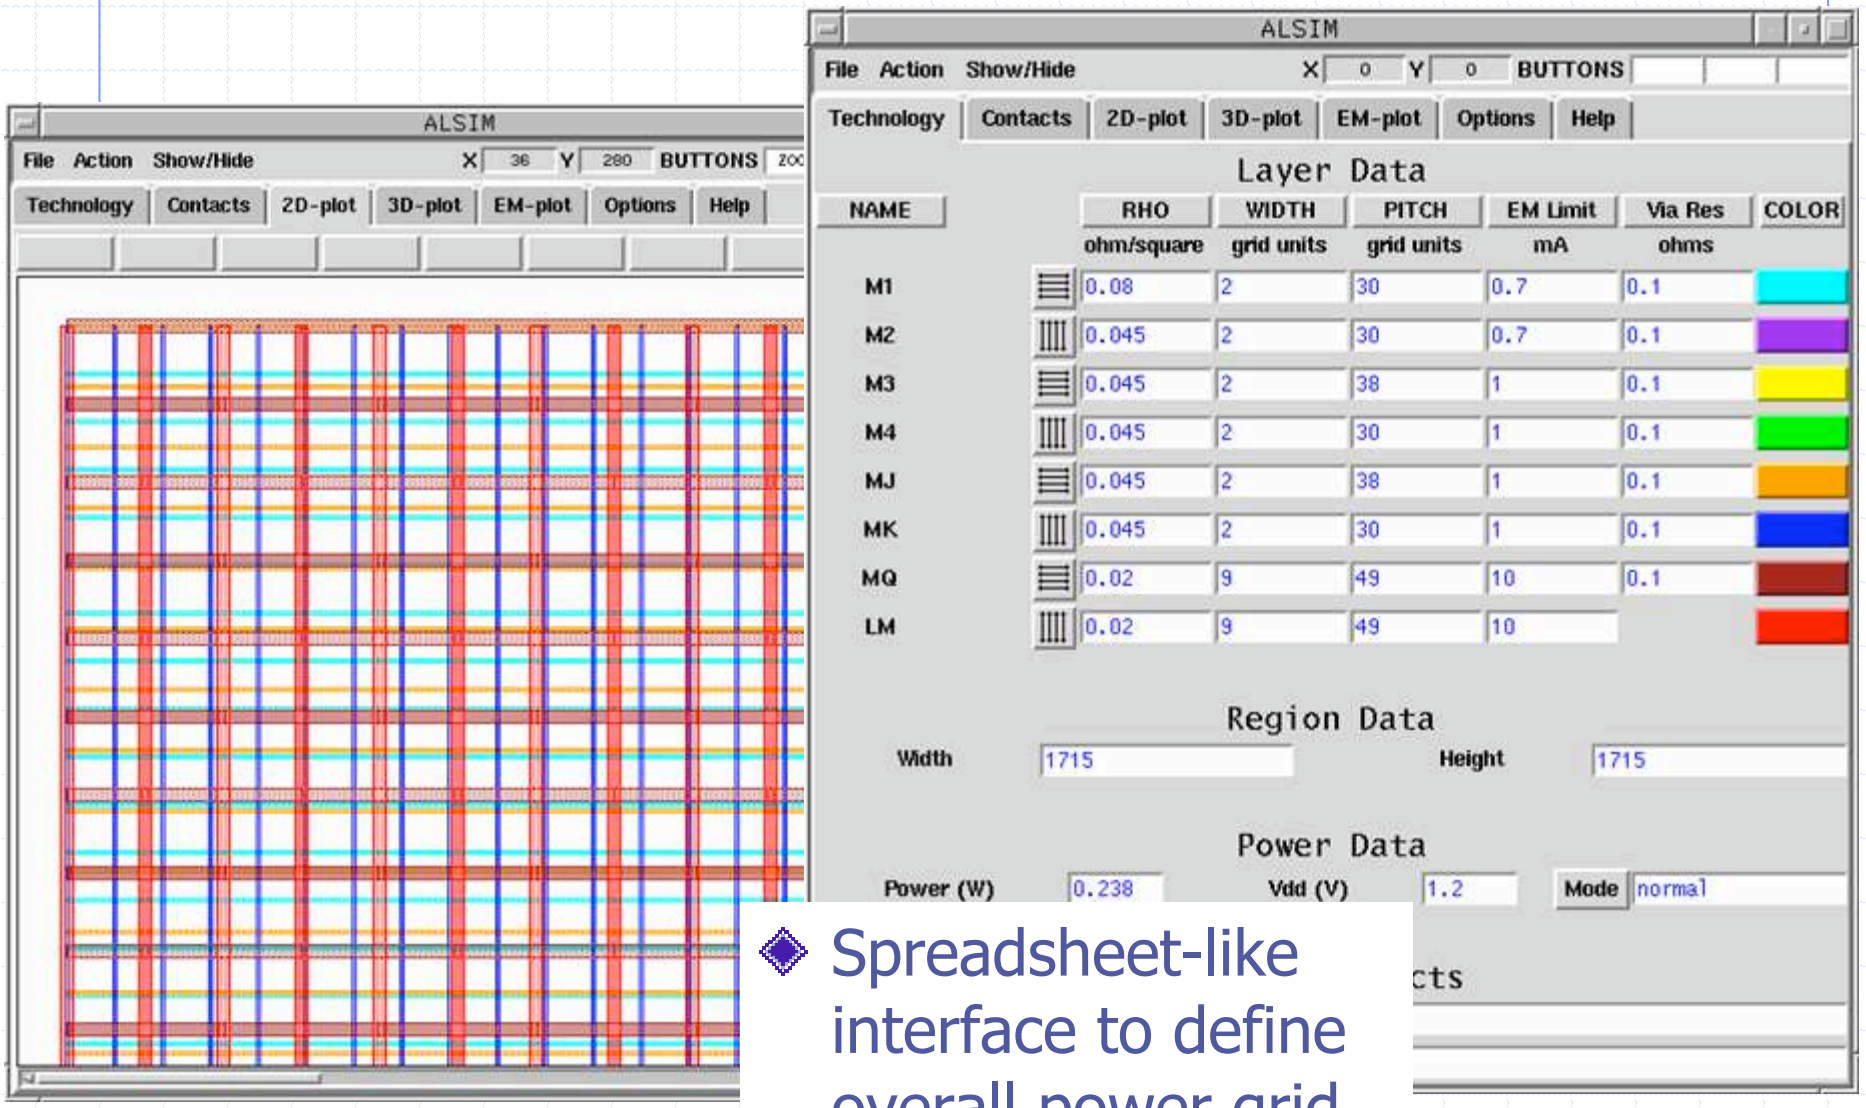

- Spreadsheet-like interface to define overall power grid.

IBM Power Grid Planner

- Lots of Visualization and Analysis…

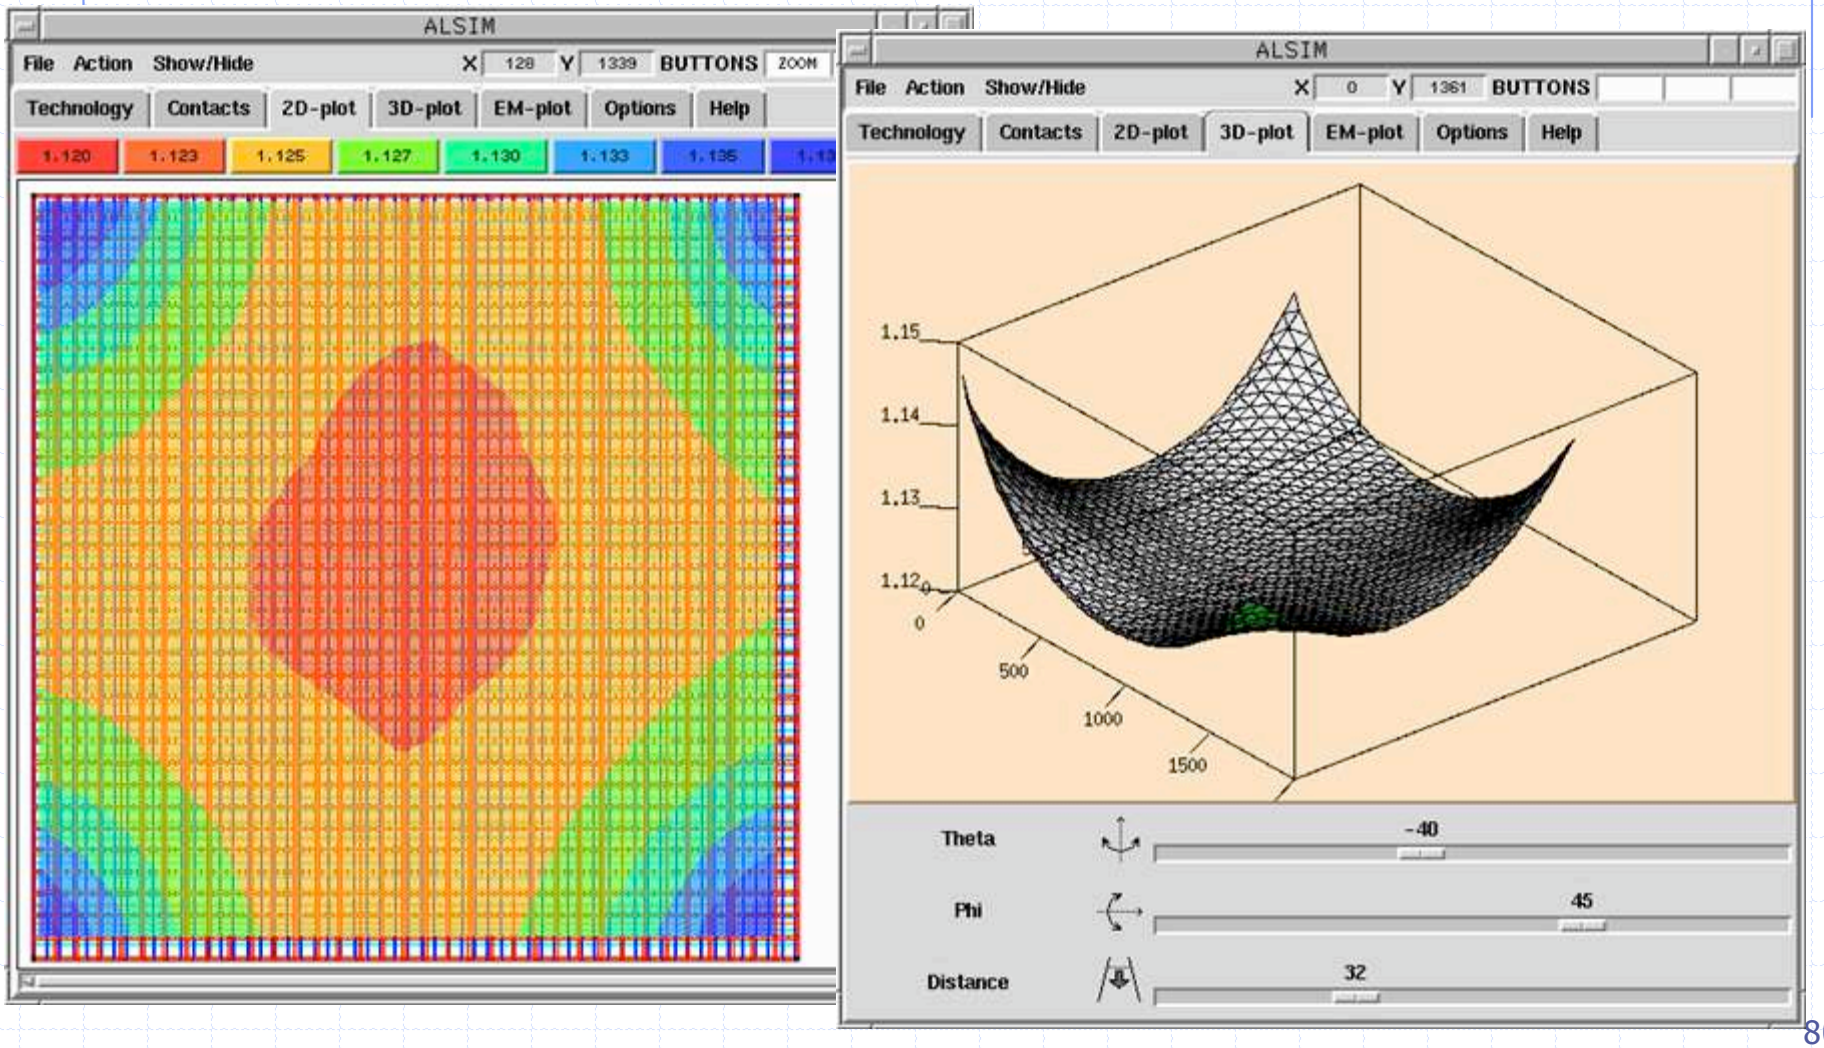

Other Tool Requirements

- As mentioned previously, automated full package analysis is a necessity!
  - Current tools require large amounts of engineering to produce results.
- Decoupling capacitance modeling is not recognized as a first-class problem.
  - Yet decoupling matters a lot for L di/dt.

Outline

- Introduction
- Variability & Uncertainty
- Power Delivery Components
  - Circuits (power dissipation)
  - Decoupling Capacitance
  - On-Chip Power Grid
  - Package
- Tool Requirements
- Conclusions

This Talk

- Defined variability and uncertainty.
- Delineated components of the power delivery system.
- Illustrated several techniques for analysis and estimation of power delivery component performance and variability.

Conclusions

- Power grid noise and its variability depends on both:
  - Technology factors.
  - Design specifics.
- Recent design trends result in a need to:
  - Perform early power delivery design.
    - Packaging technology and package selection.
    - Density and Distribution of routing resources.
    - Decoupling capacitance allocation.

Some Open Issues

- Coupling of power and timing analysis.
- Placement (PD), clock design, and power delivery analysis integration.
- Chip/package analysis and interaction.
- Vector-less Chip-level power estimation.
  - Design flows need early power analysis and decoupling estimation steps.
- Coupling of power and thermal analysis to improve reliability estimation.