Horseshoes, Hand Grenades, and Timing Signoff: When Getting Close is 'Good Enough'

Arvind NV, Krishna Panda, Anthony Hill
Texas Instruments Inc.

March 2014

Outline

- Motivation
- Uncertainty in SOC Design
- Leveraging Uncertainty
- Conclusion

MOTIVATION

Design Scaling

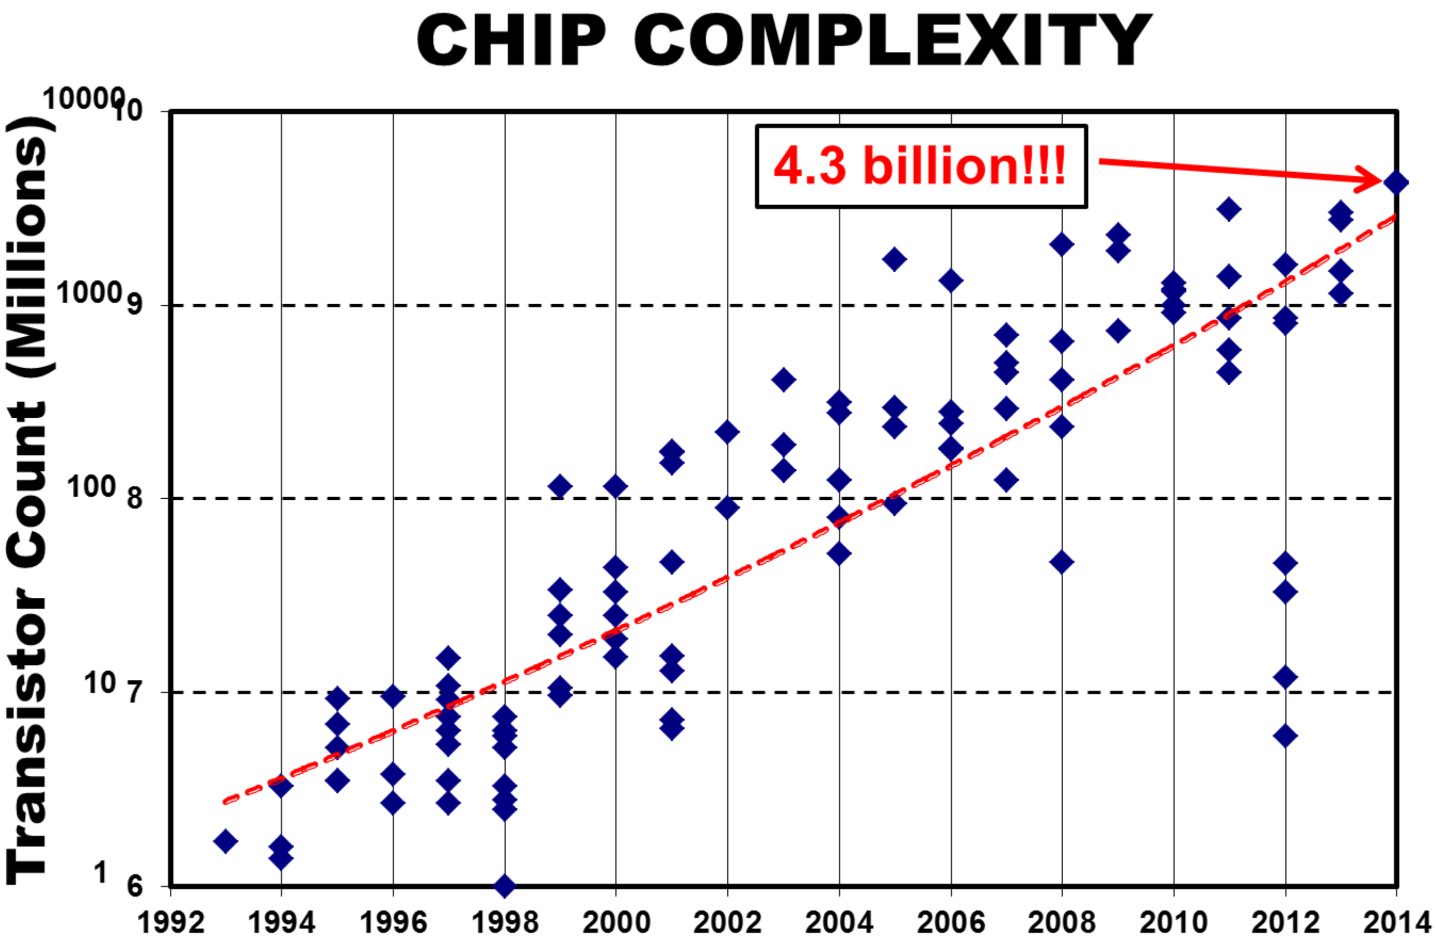

Ref: ISSCC Press Kit 2014

Tapeout Trends

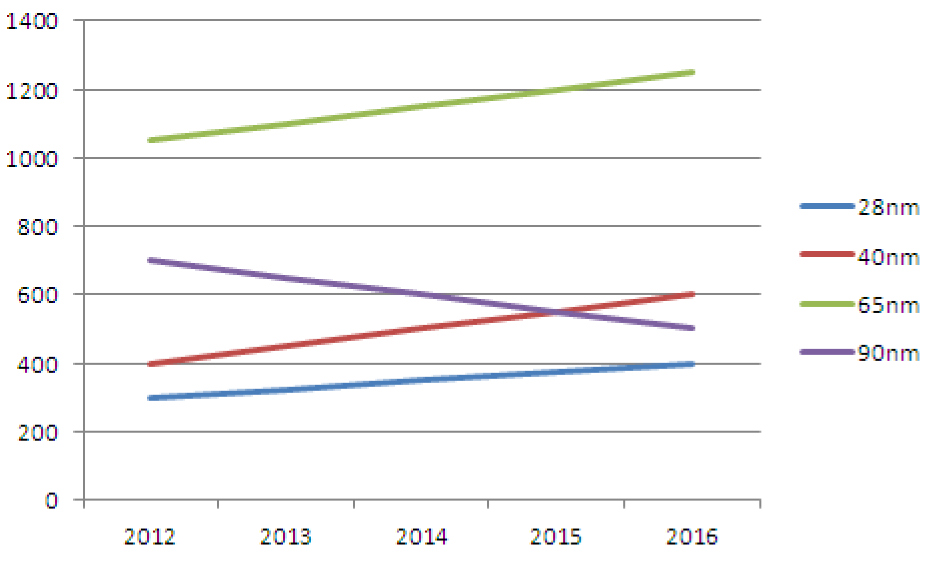

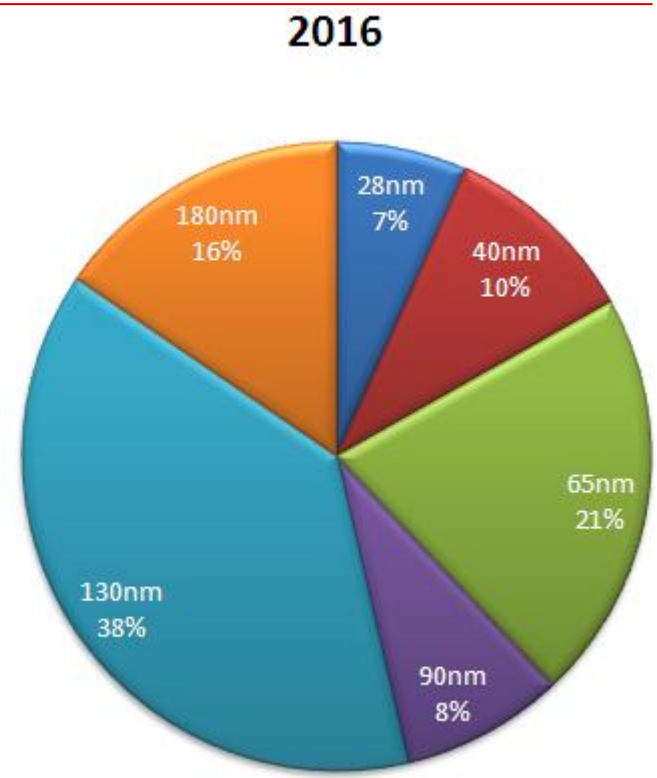

- "Mature" nodes continue to see a lot of tapeout demand.
  - In many cases, there is no benefit to advanced nodes (IO limited, cost-limited)

Ref: http://anysilicon.com/semiconductor-technology-nodes/

Scenario Complexity Circa 2006

Transistor	Voltage	Temp	RC	Family	modes				analysis			
					atpg shift	atpg capt Tclk	atpg capt Fclk	mission	setup	setup +si	hold	hold +si
Fast	Ultra-High Test	Room	minc	VBOX								
			minr									
Slow	Ultra-Low Test	Room	minc									
		Burnin	maxc									
Slow	High Burnin	Burnin	maxc	Burnin								
Fast	High Burnin	Burnin	minc									
			minr									
Slow (EOL)	Vdd - 10%	Cold	maxc	QC - MAX								
			nomc									
Slow (EOL)	Vdd - 10%	High	maxc									
			nomc									
			maxr									
Fast	Vdd + 10%	Cold	minc	QC - MIN								
			minr									
Fast	Vdd + 10%	Hold	maxr									
Typical	Vdd	Room	nomc	Typical								
"Near Fast"	Vdd - AVS	High	maxc	AVS								

- This has increased significantly with widespread adoption of AVS and DVFS.

AVS & DVFS

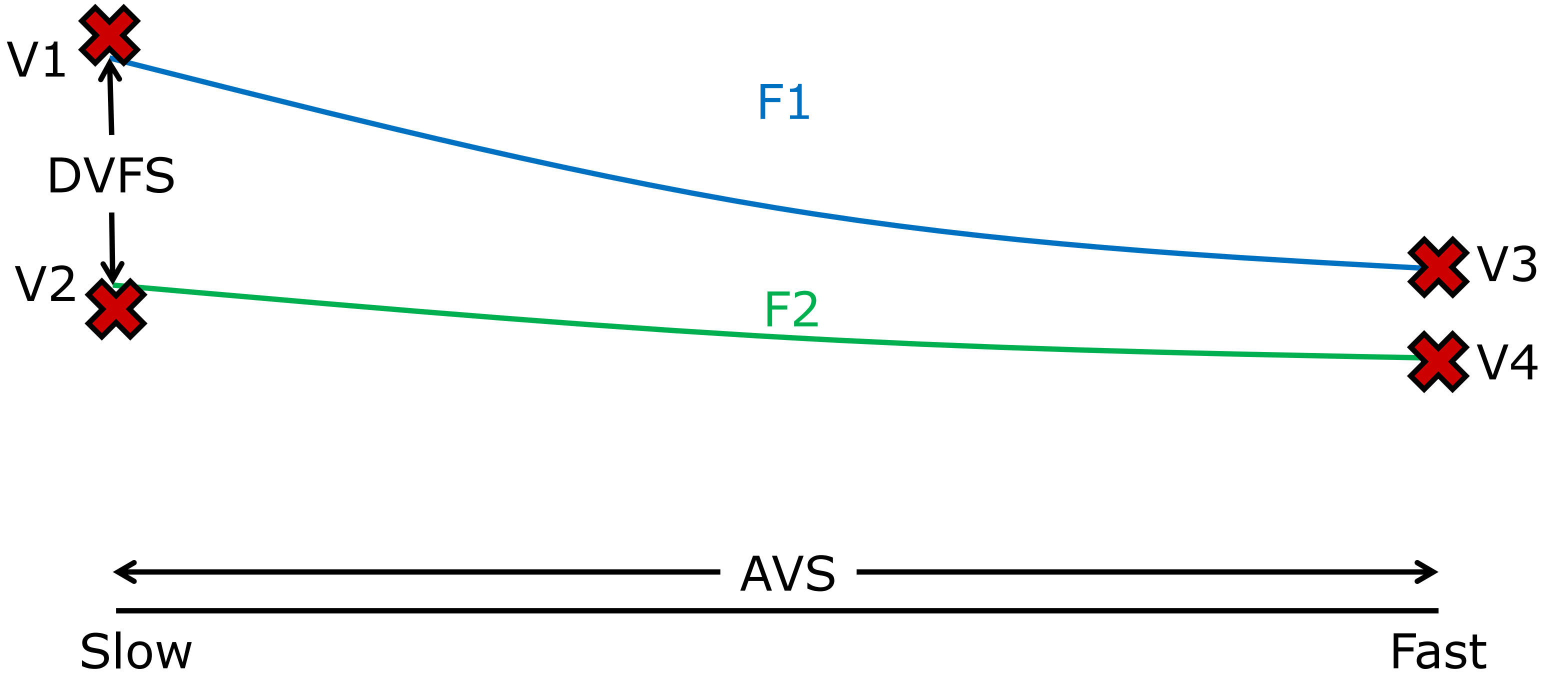

- Voltage scaling – to reduce power at lower frequencies or to reduce power for fast process corners – has increased the risk of 'outliers' and hence, the need to analyze additional PVT scenarios.

Example: Silicon Prediction

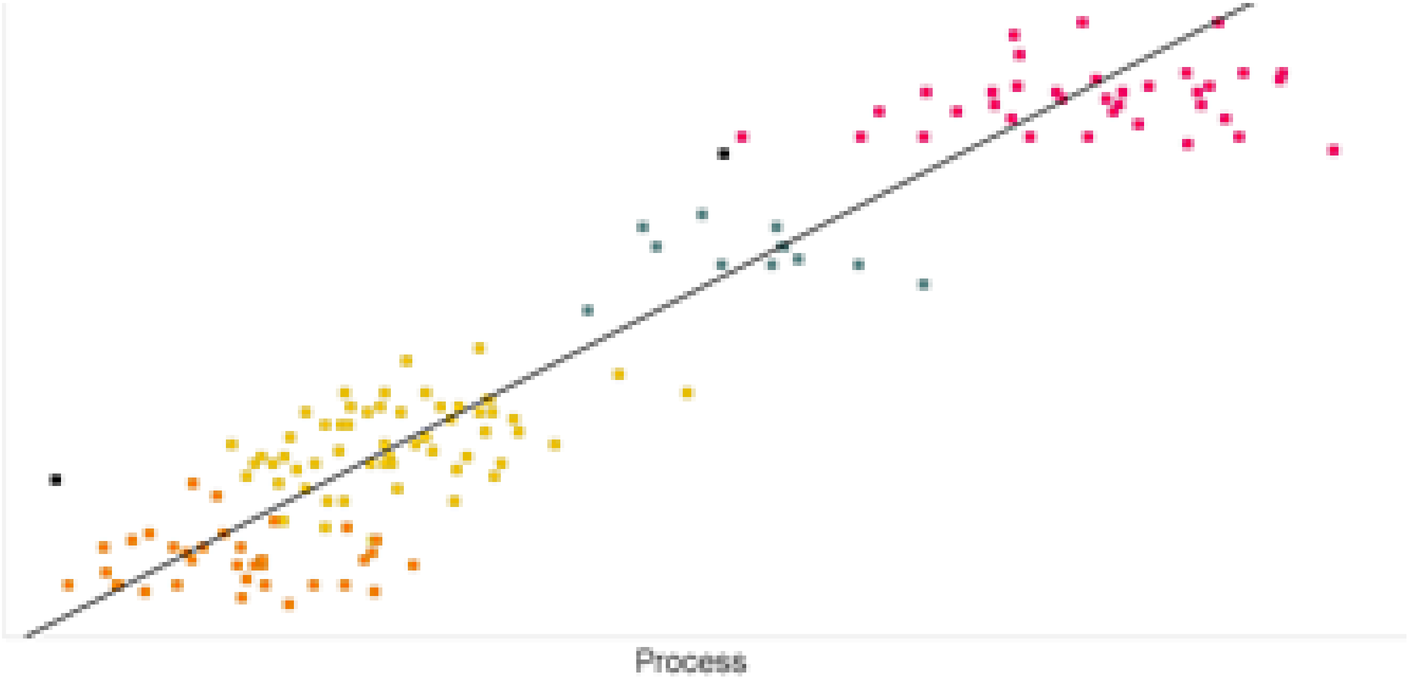

UNCERTAINTY IN SOC DESIGN

Local Mismatch

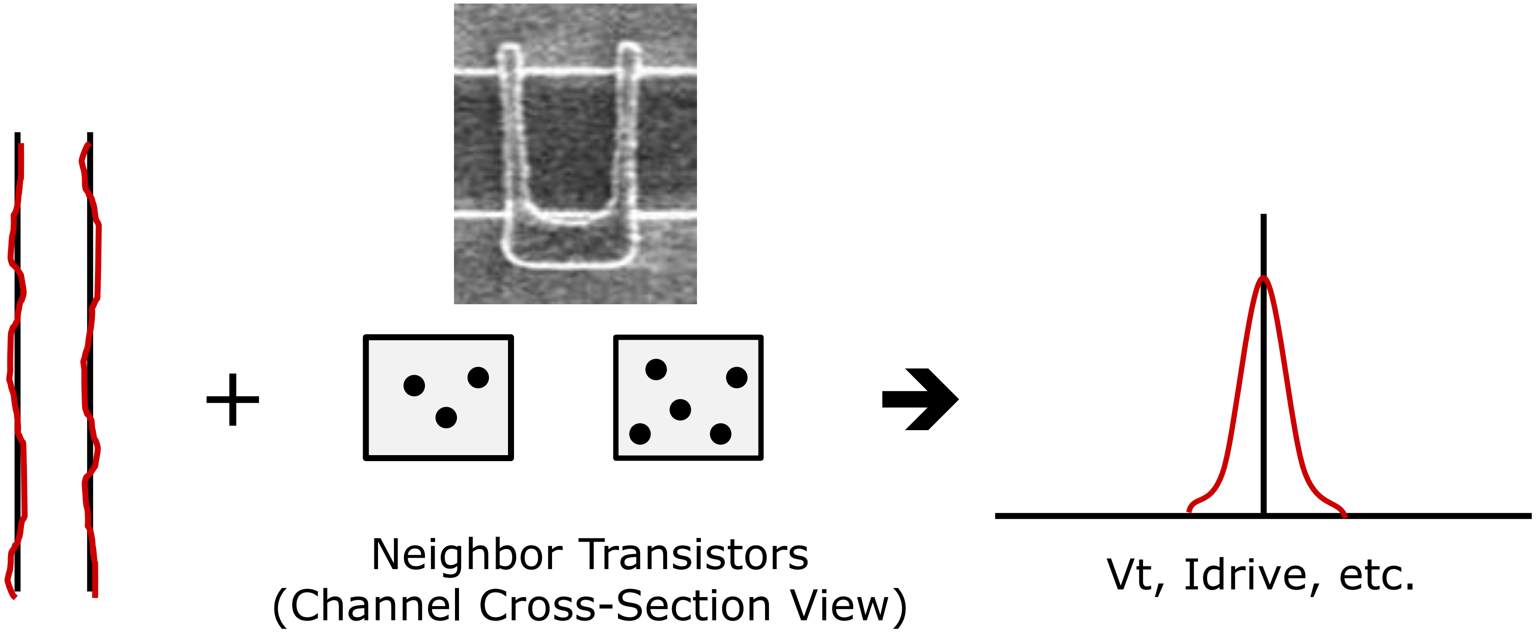

- Performance of neighboring transistors don’t match.
  - Line edge roughness (LER): no edges are perfectly straight.
  - Random dopant fluctuation (RDF): channels have varying dopants.
  - These effects (and others) create **<u>local mismatch</u>**.
- Local mismatch is generally increasing node-to-node.
  - SPICE models typically account for some (not all) local mismatch.

SPICE Model "Uncertainty"

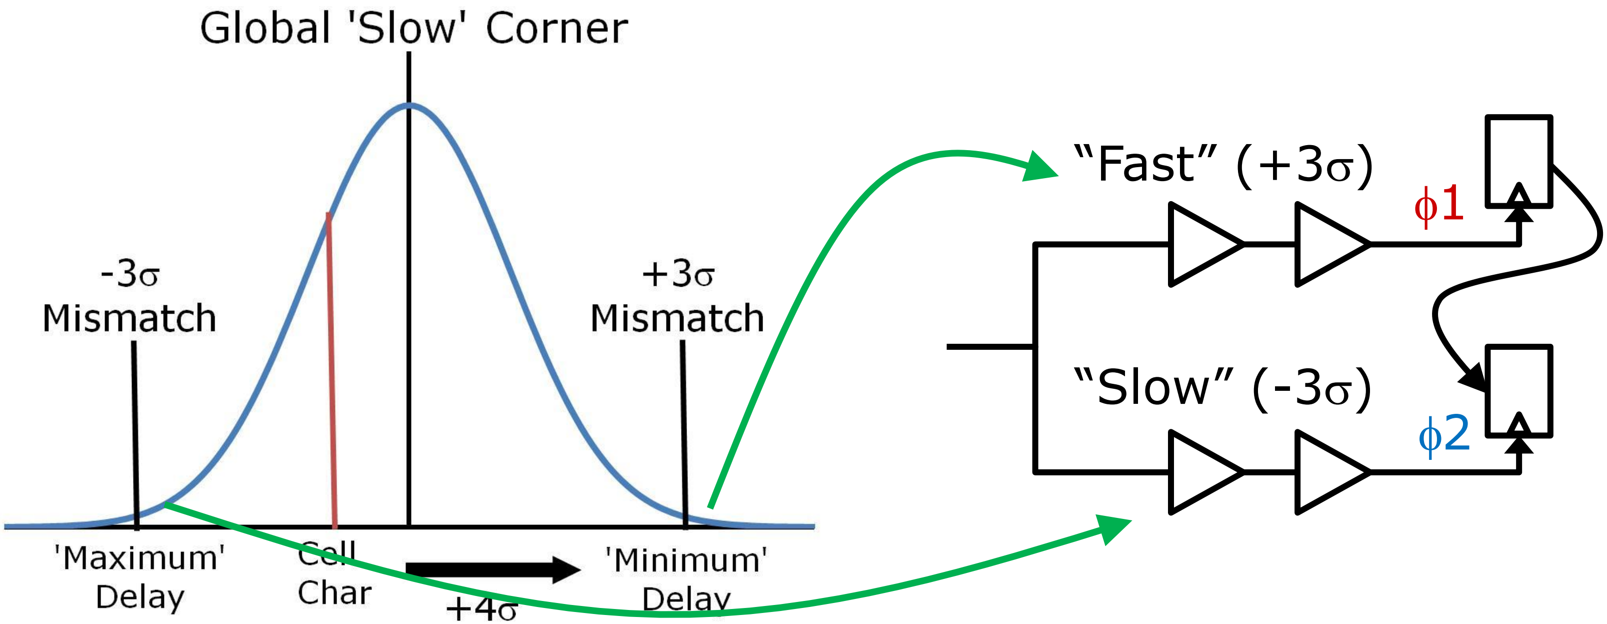

- "Corner" models are not bounding.
  - Differential delay (race) conditions exist on an SOC.
  - E.g., launch and capture clocks for hold-time checks
- What is in your timing characterization?
  - If pessimistic for small cells, how much faster are large cells?

28nm Local Mismatch (SiON)

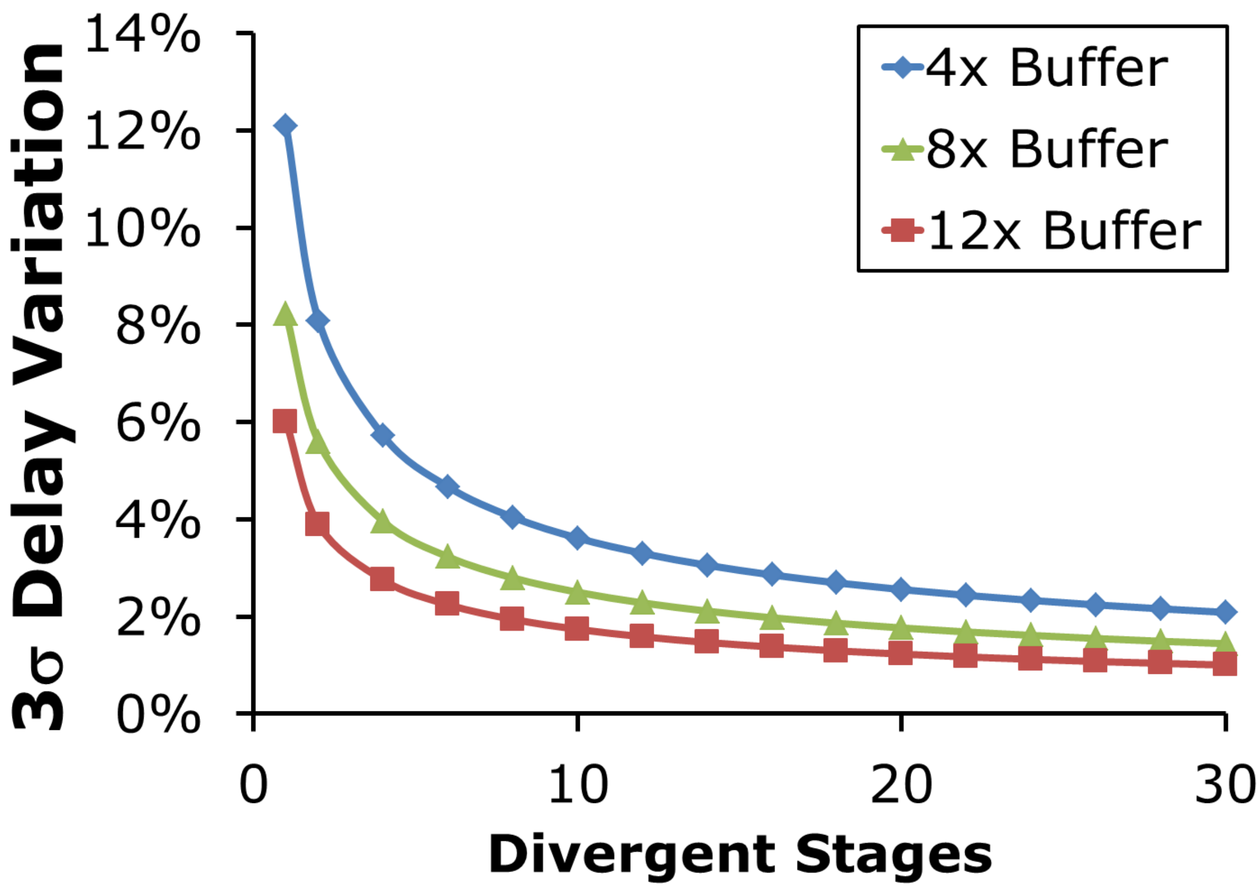

Cell Context Variation

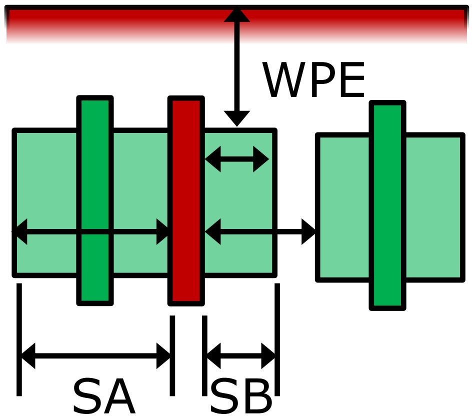

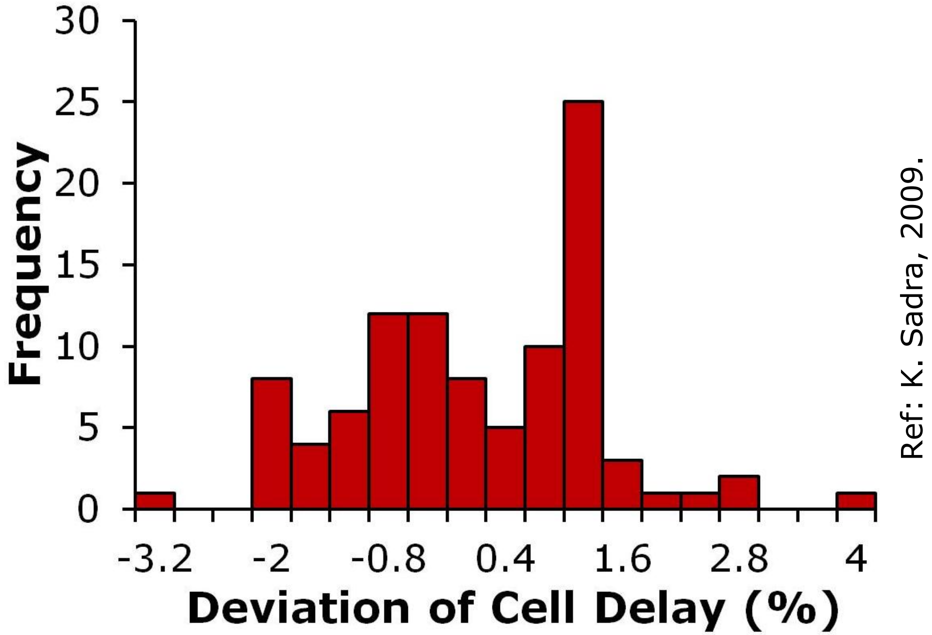

- Cell performance depends on its environment.
  - Gate distance to diffusion edges – Length of Diffusion (LOD)
  - Gate distance to well edges – Well Proximity Effect (WPE)
- I_{drive} can vary by 10-20% (more if not managed properly).

Dynamic IR Drop

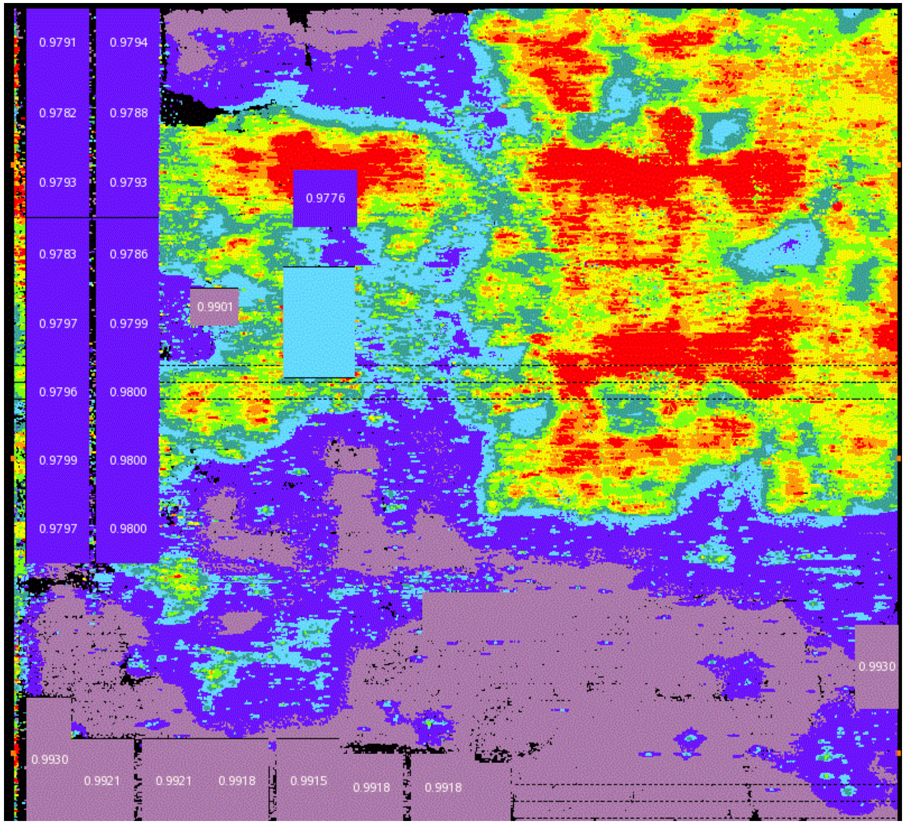

- Dynamic IR drop can change significantly across even small distances on an SOC.
  - Different clock domains, logic depth, decoupling cap density.

Dynamic IR

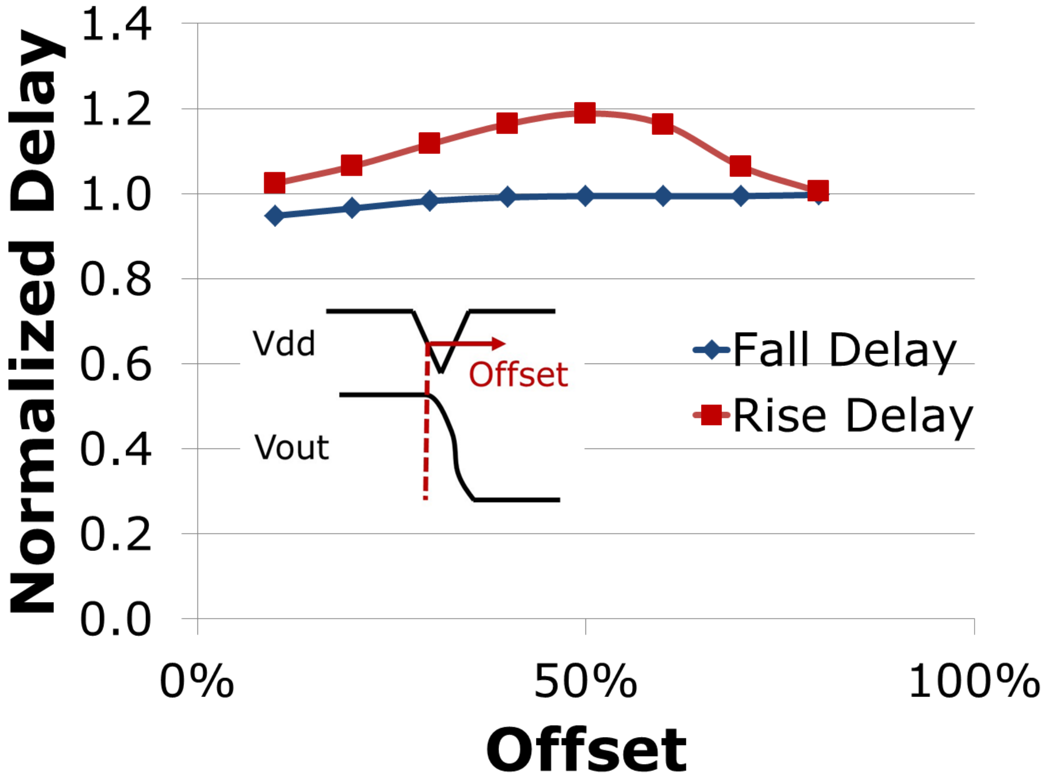

- Dynamic IR can speed up or slow down logic gates.

Parasitic Accuracy

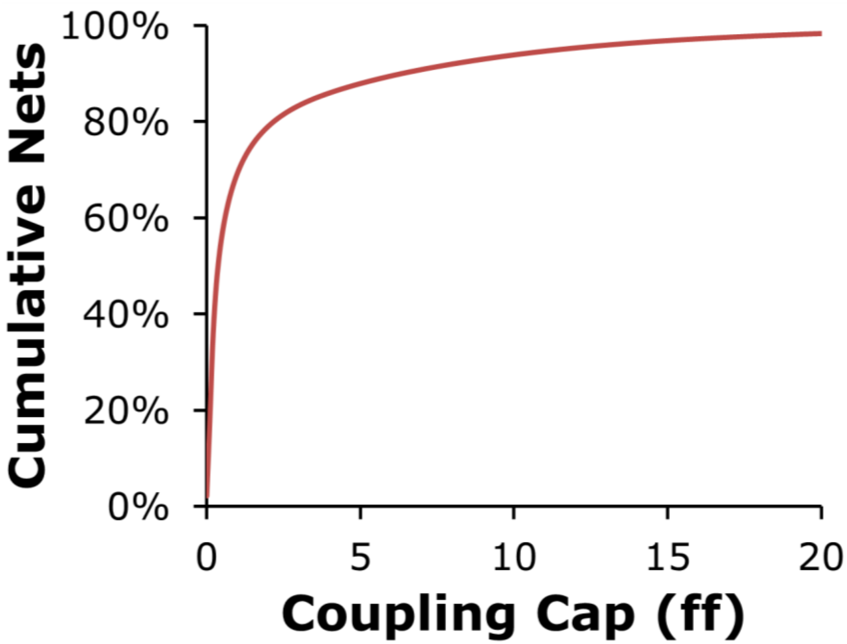

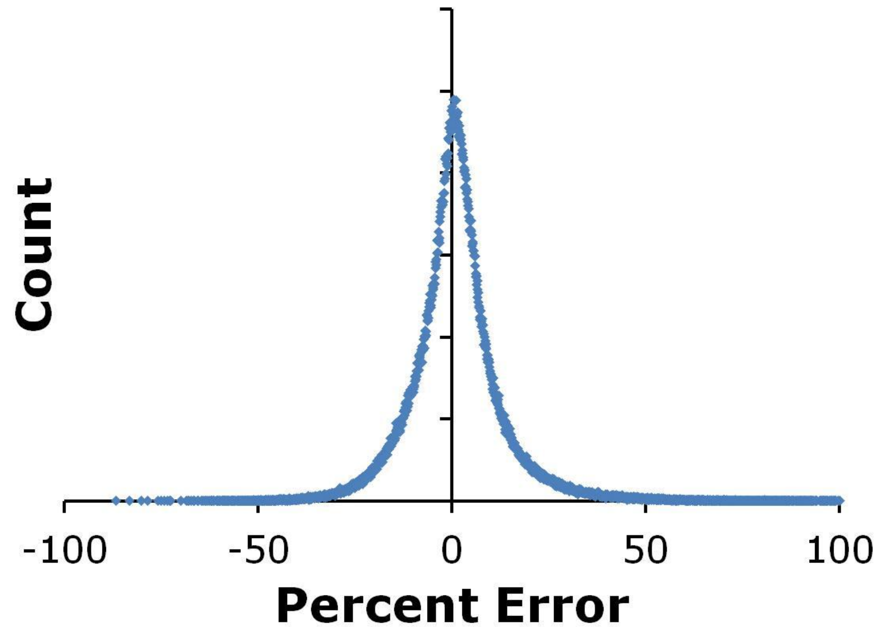

- The majority of wire-to-wire coupling involves small capacitances.
- At 28nm, >80% of net-to-net coupling is <5ff.
- The large number of SOC geometries and run time limit our ability to deploy true 3D simulation for capacitance.
- The net result is that error on these caps is typically 20-100%.

Inter-Layer Metal Mismatch

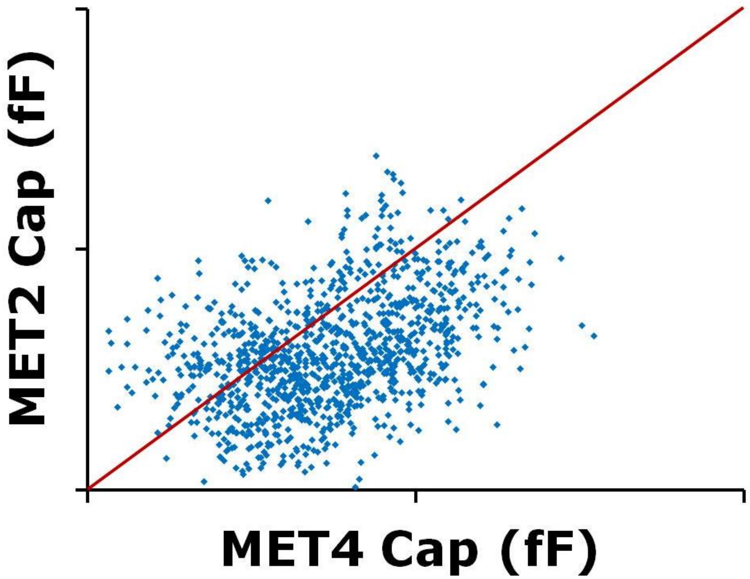

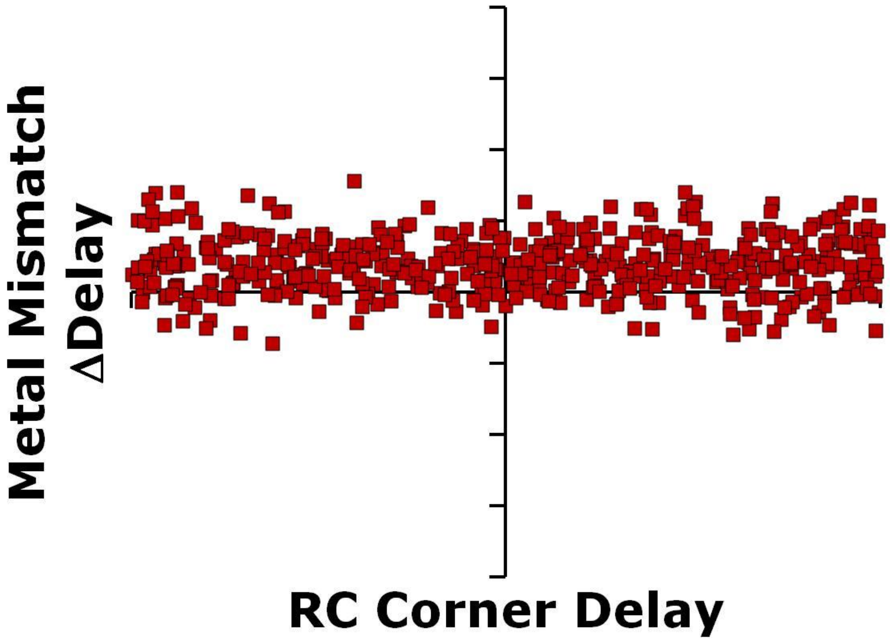

- PTV scenarios assume a specific interconnect with matched layers.
  - A corner assumes all layers are at one single condition (e.g., cbest).
  - In reality, each layer is constructed independently and may vary.
  - E.g., M3 may have max etch, M4 may have minimum etch.

Multi-Vt Process Skew

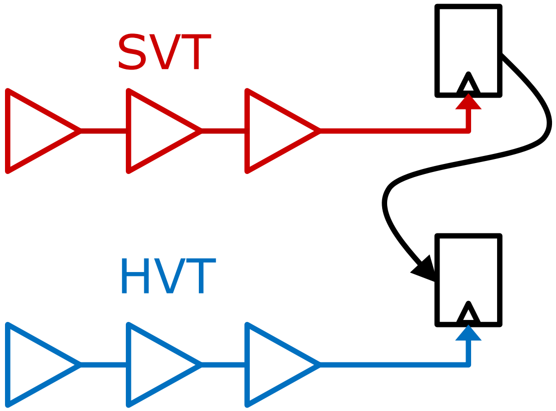

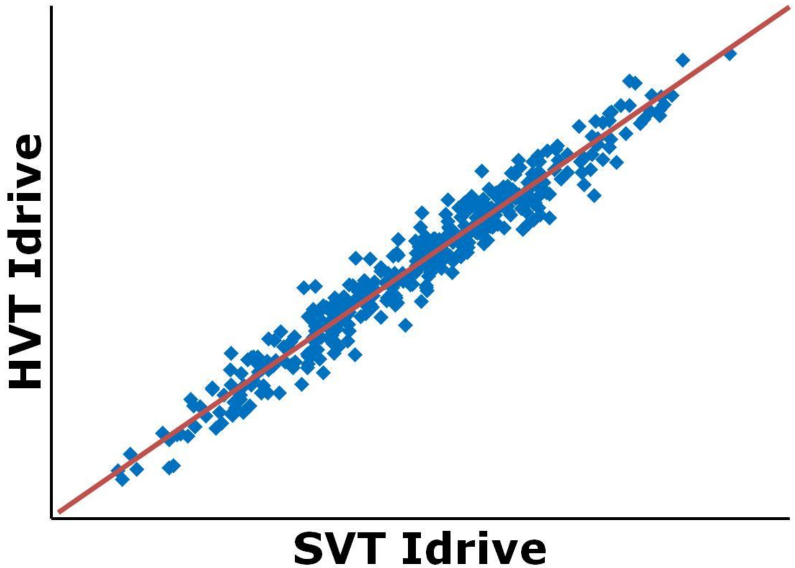

- Devices with different Vt targets are not precisely correlated.
  - Implants tend to be independent.
  - E.g., design may be closed with SVT and HVT both at the fast corner, but hold fallout occurs when HVT runs slightly 'colder'.
- Multiple Vt devices are often mixed on timing paths.
- This makes it challenging to predict actual path performance.

Aging

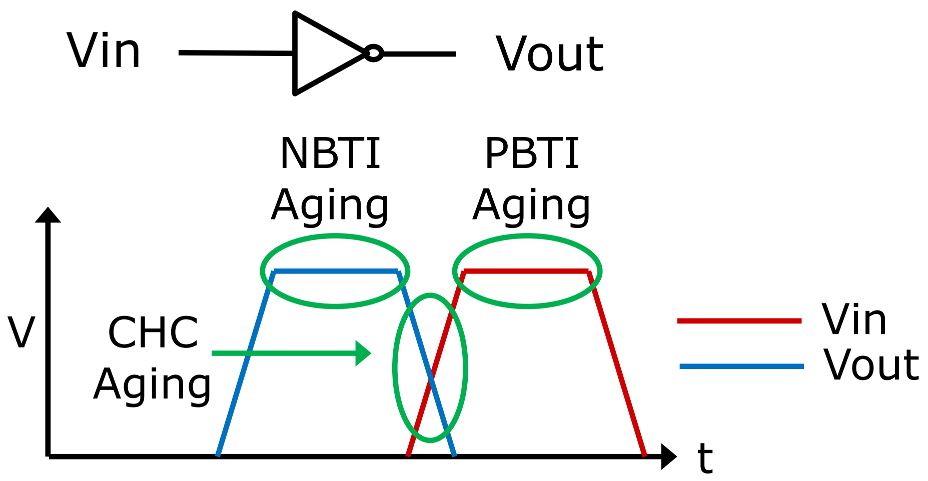

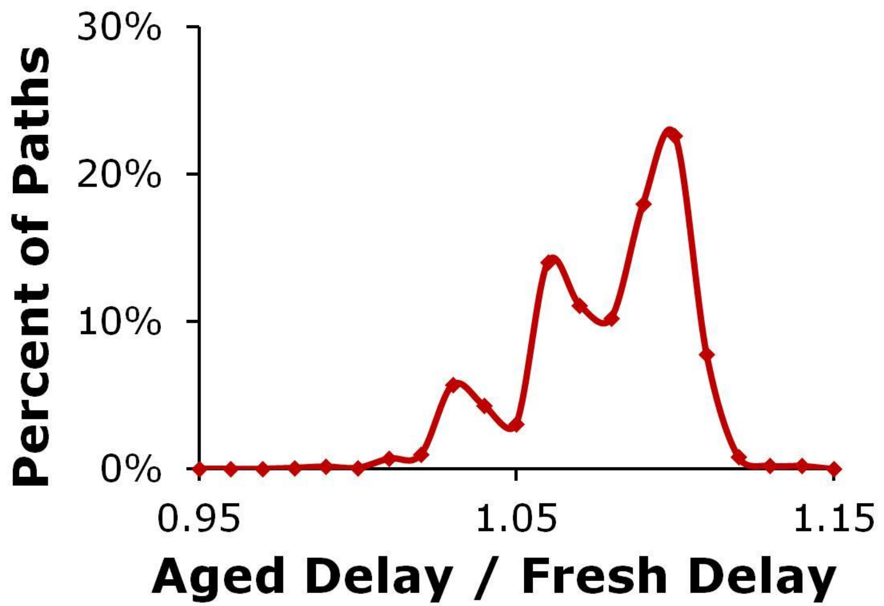

- Devices age due to gate and drain stress.
  - The net effect can be either speed up or slow down of a path.
- Implementing a block characterized with fresh timing models then timing with a library characterized at 100k PoH shows up to a 15% timing degradation.

Clock Aging

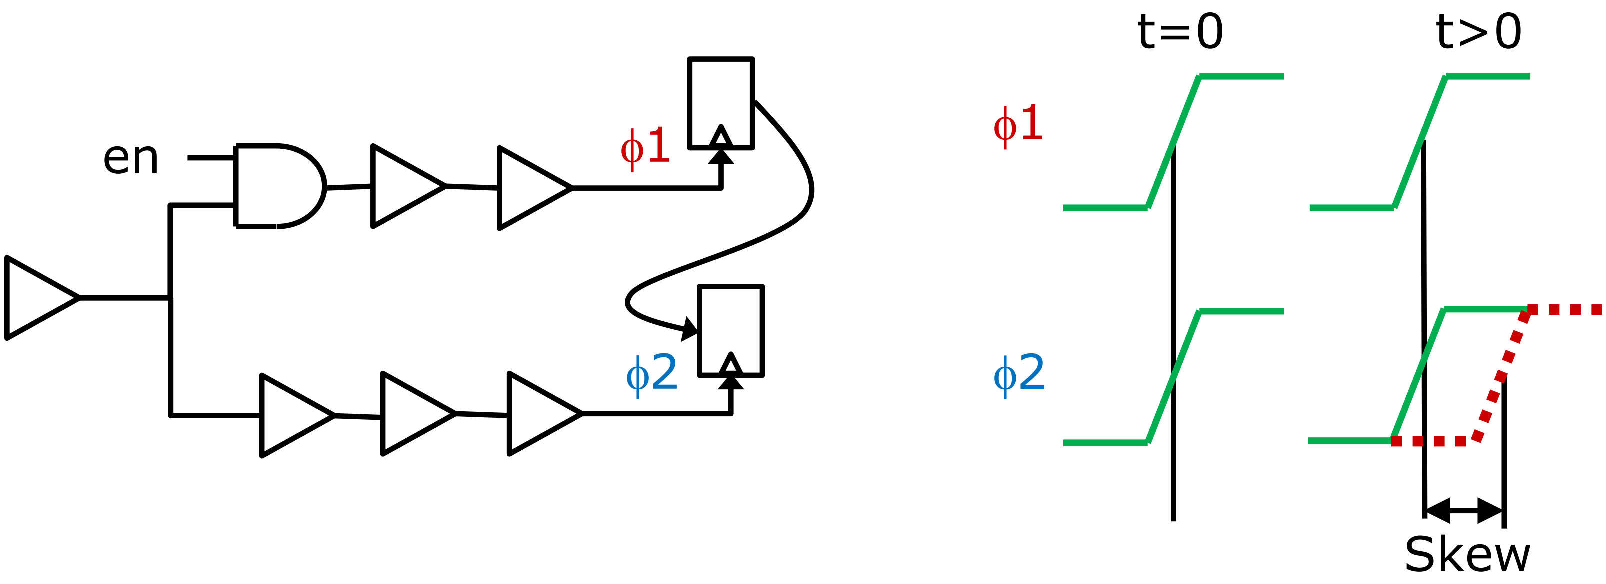

- Clock gating is a very common methodology in SOC design.
- Gating clocks creates age-based skew in the clock tree.
  - Aged skew can be huge – (100ps+ for deeply-gated trees).
- The amount of aging varies based on a history of how often the clocks are gated.

Other Uncertainties

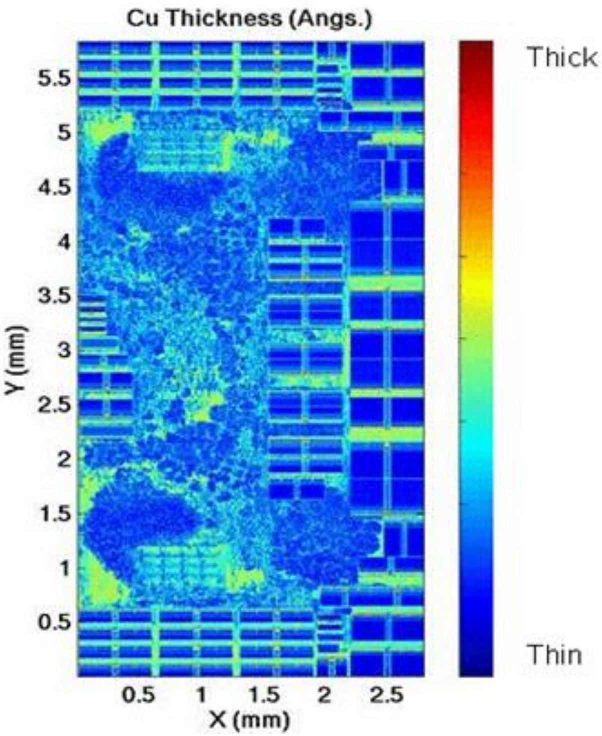

Metal Thickness

Temperature

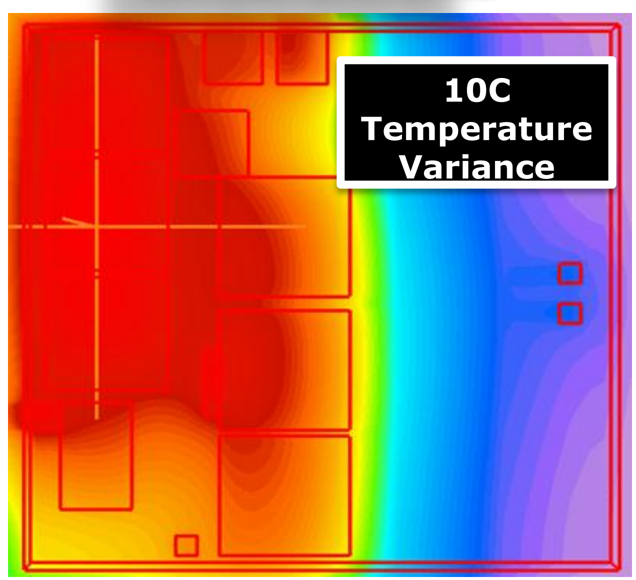

Range	No. of Paths (Old)	No. of Paths (New)
-2% to -1%	0	2
-1% to 0%	0	39
0% to 1%	34	114
1% to 2%	175	370
2% to 3%	180	757
3% to 4%	471	471
4% to 5%	349	54

STA Engine 'Errors'

LEVERAGING UNCERTAINTY

Time-to-Tapeout

- Understanding the uncertainty in design can be used to improve time-to-tapeout.

- Fewer ECO Loops
  - e.g., through better implementation-to-signoff correlation
- Run-Time
  - e.g., reduced parasitics, simpler timing models
- Memory
  - e.g., reduced parasitics
- Compute
  - e.g., fewer scenarios

Coupling: Small Aggressors

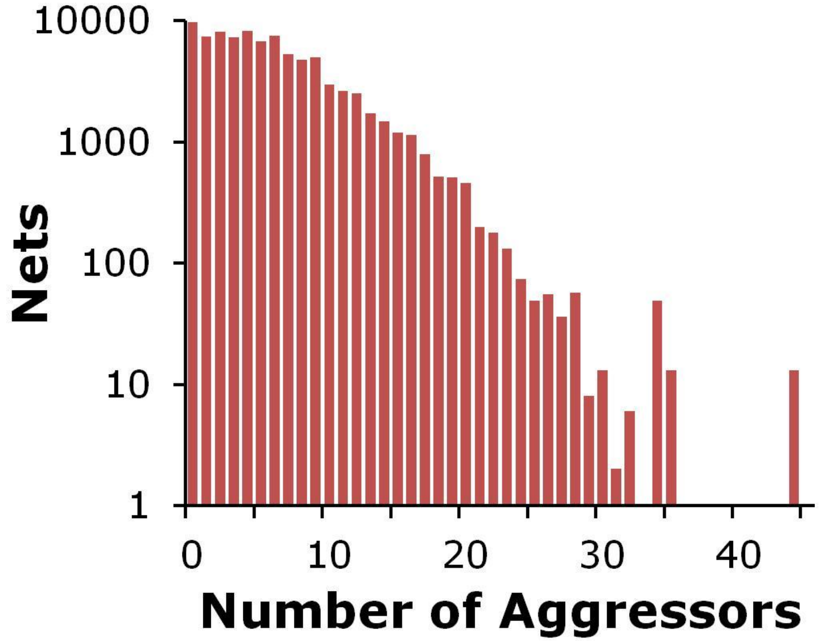

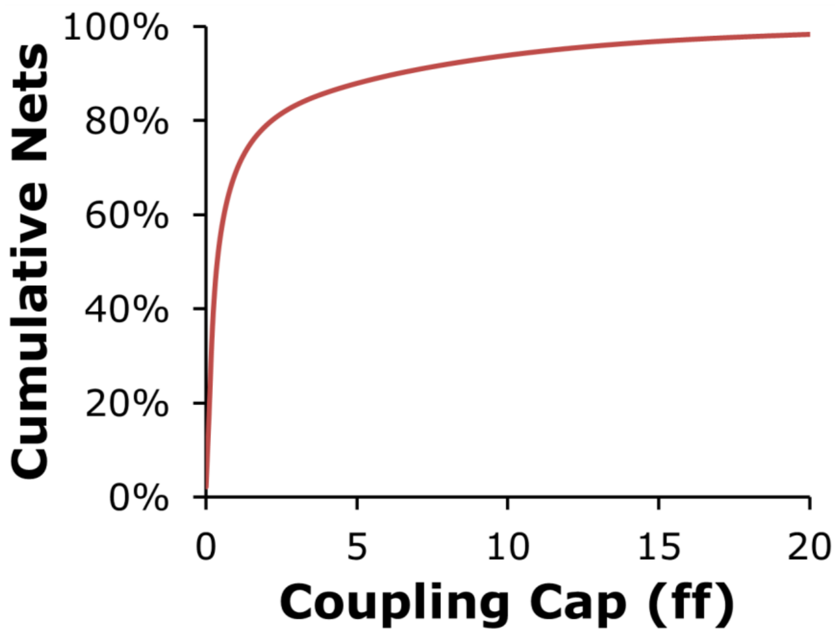

- Most aggressor-victim pairs have tiny coupling capacitance.
  - (And there is high inaccuracy on these small coupling caps.)
- We can improve the "SI Experience" by intelligent filtering.
  - Filter based on aggressor / victim relationships
  - Grouping small aggressors
  - Ignoring small aggressors

Small Aggressor Modeling

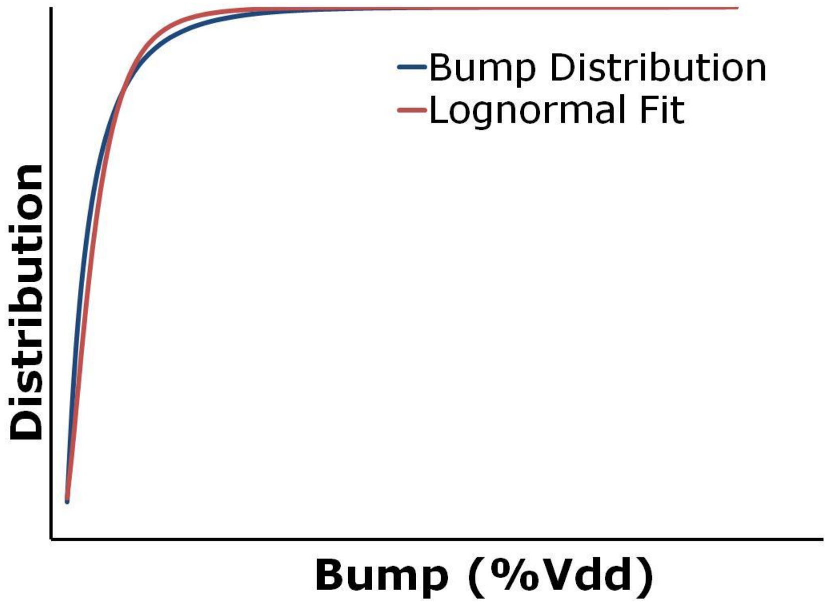

Probability of Timing Escape						
Path Error	Aggressor Bump					
	0.5%	1%	2%	3%	5%	10%
0.1%	59%	77%	88%	92%	95%	97%
0.5%	7%	27%	52%	65%	77%	88%
1.0%	1%	7%	27%	42%	59%	77%
2.0%	0%	1%	7%	18%	35%	59%
3.0%	0%	0%	2%	7%	21%	46%
5.0%	0%	0%	0%	1%	7%	27%

Error on a 750ps Clock Cycle

- Empirically, the small-aggressor timing impact on a net can be modeled as a log-normal distribution.
- We can calculate error vs. accuracy using statistical methods.
  - With appropriate assumptions on gate delay, number of gates, …
- This provides a framework to trade-off run-time and accuracy vs. design margin and risk.

Small Aggressor Filtering

Filter Threshold	TNS	WNS	Violation Count
0.005	-13.1	-0.067	1126
0.01	-2.72	-0.049	323
0.02	-0.59	-0.047	52

- Aggressive filtering of small aggressors can pay dividends on reduced timing violations, ECOs, and time-to-tapeout.

Crosstalk on Clock Nets

- Crosstalk on clock increases timing closure effort.
  - Can be a significant source of pessimism.
- Fix outliers and then ignore (disable) crosstalk-induced delay on clock.
- This methodology has successfully been deployed across multiple technology nodes.

Crosstalk Delay (ps)	Number of Nets
0	9356
0.5	0
1	19
1.5	41
2	21
2.5	11
3	2
3.5	1

Sensitivity-Based Signoff

- Multiple scenarios across PTV and RC serve to highlight paths which have sensitivity to process or environment.

- Eliminating sensitive circuits will enable reduction of scenarios which vary only in process, temperature, voltage, or interconnect corner.

- These methods may include:
  - limiting wire length (and RC)
  - strict max cap limits
  - smart usage of small drive cells
  - limiting crosstalk (large bumps, noisy slews)
  - <u>➔ crosstalk as a DRV!</u>
  - elimination of SI-induced bumps on clock

Example: Small Cell Handling

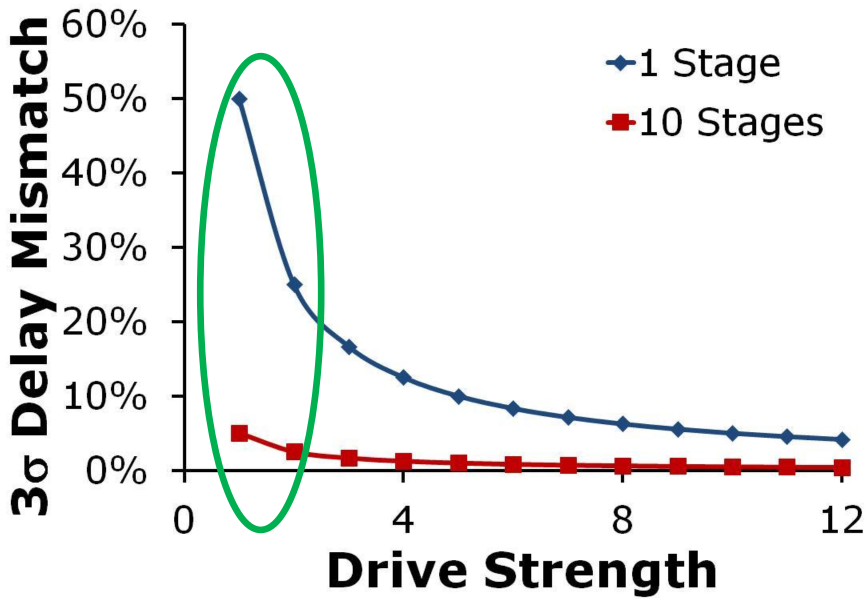

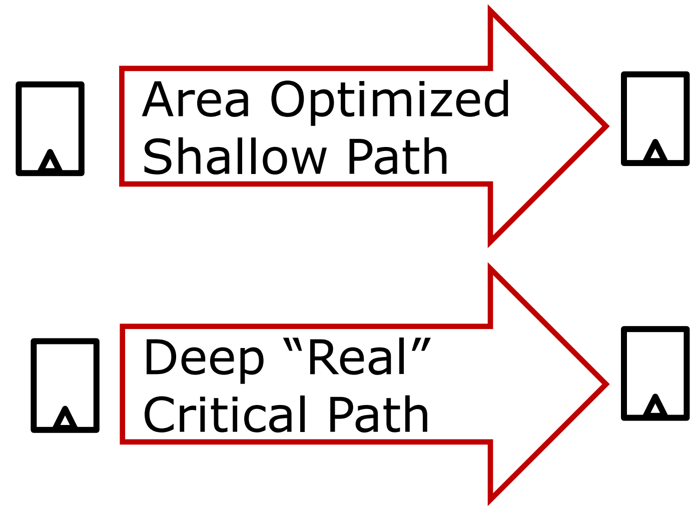

- Small transistors are highly sensitive to variation.
  - Optimization creates small-cell dominated critical paths.
- We desire to avoid small cells on near-critical timing paths.
  - Datapath depth-based derating ➔ computationally complex.
  - Post-optimization analysis + targeted fixing ➔ time intensive.
  - Derate timing on small cells ➔ practical with minimal impact.

Example: Clock Skew Sensitivity

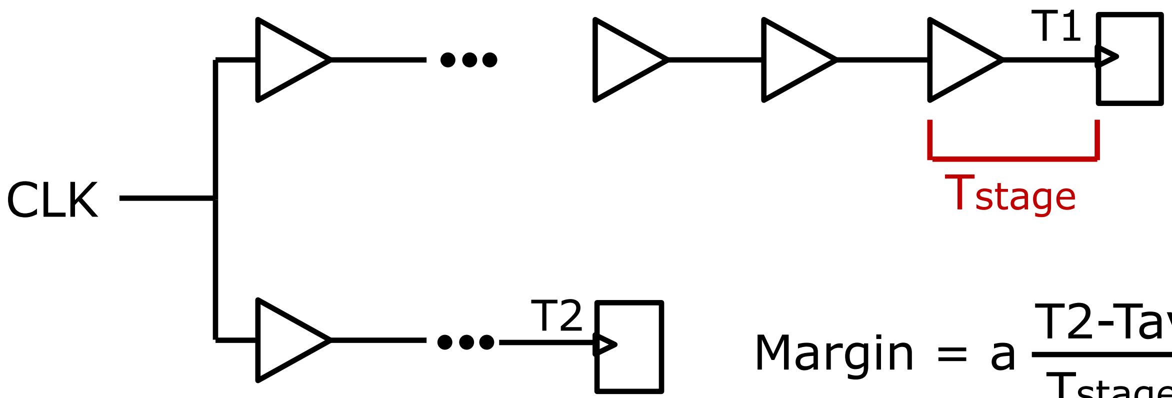

$$\text{Margin} = a\,\frac{T2 - T_{avg}}{T_{stage}}$$

- Skewed circuits often show variation across PTV.
  - Launch and capture edges do not track across all RC or gate delays.
- Targeted margins can eliminate the need to analyze this.
  - Any RC or gate mismatch between launch and capture are covered by a margin.

CONCLUSIONS

"Close is Good Enough"

- STA prediction of silicon performance is generally poor.
  - Unknowns permeate SOC design: characterization, coupling, model accuracy, on-die variation, metal mismatch, etc.

- Understanding these uncertainties can reduce complexity in STA signoff and speed time-to-tapeout.

- <u>Sensitivity-based signoff</u> would significantly reduce signoff scenarios.
  - e.g., DRV checks for wire length, RC, max SI bump/delay, and max noisy slew have been proposed to reduce outliers.