

Predictive Technology for Advanced Node Design Exploration

Rob Aitken
ARM R&D



What's Ahead?

The Scaling Wall?

EUV around the corner?

Slope of multiple patterning

Avalanches from resistance, variability, reliability, yield, etc.

Crevasses of Doom

Scaling getting rough


Trade off area, speed, power, and (increasingly) cost

A Digression on Node Names

- Process names once referred to half metal pitch and/or gate length
 - Drawn gate length matched the node name
 - Physical gate length shrunk faster
 - Then it stopped shrinking
- Observation: There is nothing in a 20nm process that measures 20nm

Node	1X Metal Pitch
Intel 32nm	112.5nm
Foundry 28nm	90nm
Intel 22nm	80nm
Foundry 20nm	64nm

Sources: IEDM, EE Times




14nm node ARM M1 clip, 46nm minimum pitch, exposed on an NXE:3300B with conventional illumination

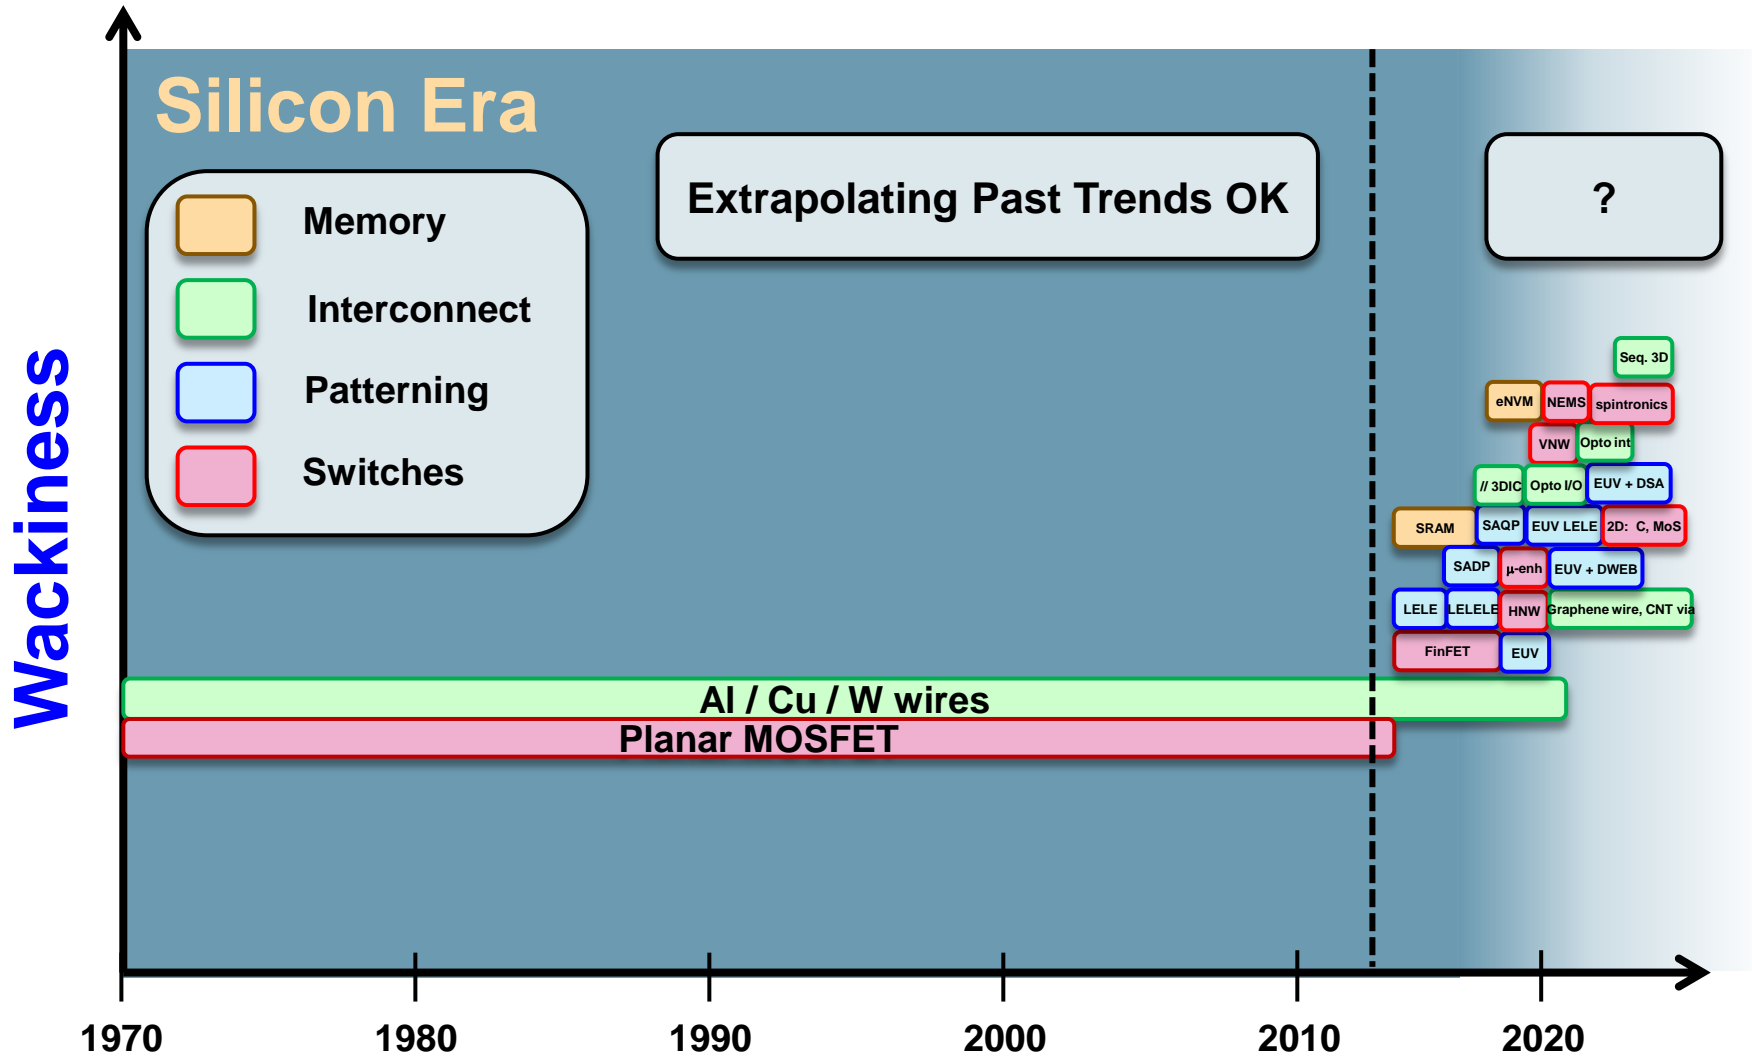
Source: ASML keynote, IEDM 12



“Silicon” Device Roadmap



“Silicon” Device Roadmap



Where is this all going?

- Direction 1: Scaling (“Moore”)
 - Keep pushing ahead $10 > 7 > 5 > ?$
 - N+2 always looks feasible, N+3 always looks very challenging
 - It all has to stop, but when?
- Direction 2: Complexity (“More than Moore”)
 - 3D devices
 - eNVM
- Direction 3: Cost (Reality Check)
 - Economies of scale
 - Waiting may save money, except at huge volumes
 - Opportunity for backfill (e.g. DDC, FDSOI)
 - For IOT, moving to lower nodes is unlikely
- Direction 4: Wacky axis
 - Plastics, printed electronics, crazy devices



Three Questions

1. How fast will a CPU go at process node X?
2. How much power will it use?
3. How big will it be?

**How close are we to the edge?
Is the edge stable?**



The Answers

1. How fast will a CPU go at process node X?
 - Simple device/NAND/NOR models overpredict
2. How much power will it use?
 - Dynamic power? Can scale capacitance, voltage reasonably well
 - Leakage? More than we'd like, but somewhat predictable
3. How big will it be?
 - This one is easiest. Just need to guess layout rules, pin access and placement density. How hard can that be?

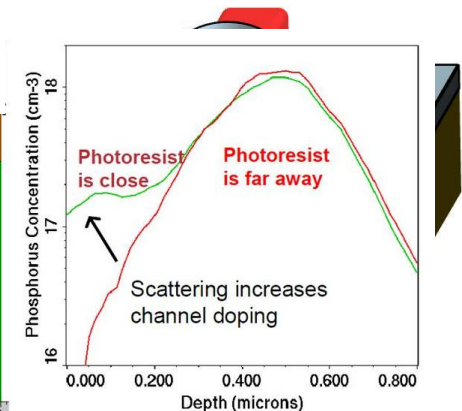
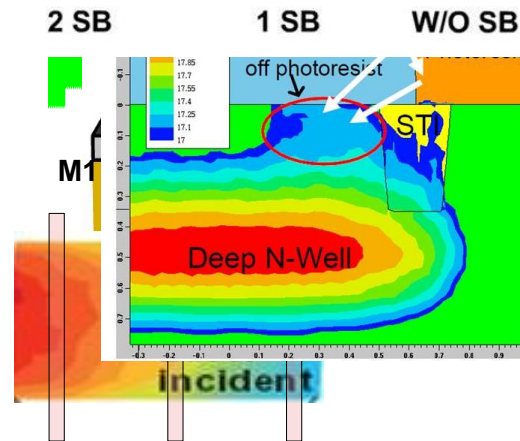

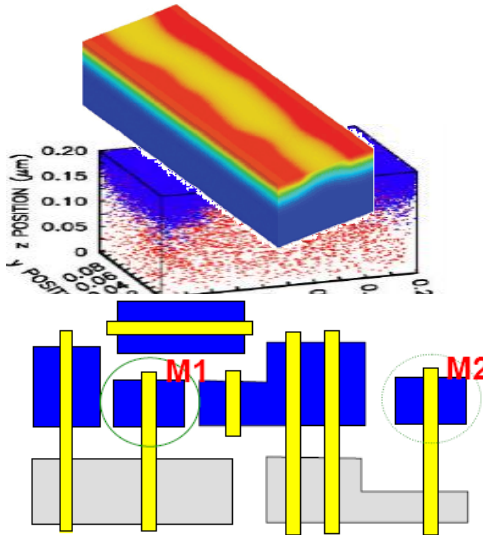
DTCO Over Time

- \geq 90nm: DFM rules, customized from recommendations
- 65nm: Quantified DFM Metrics
- 45/40: Regularization, restricted pattern sets
- 32/28: Early discussion, rule/device suggestions, early test vehicle
- 22/20: Joint development of rules & process, double patterning, predictive technology
- 16/14: Improved structures, CPU demonstrators, collaboration with equipment vendors
- 10nm: Stay tuned – still early. Silicon validation

Cell Library Design Context Issues


- Random variability
 - LER
 - RDF
- Systematic variability (context)
 - OPC interactions
 - Lithographic fidelity
 - Rounding
 - Strain engineering
 - Well Proximity Effect
 - Millisecond Anneal

900+ standard cells per library
(V_T , L, drive strength, etc.)





CPU Core Design Context Issues

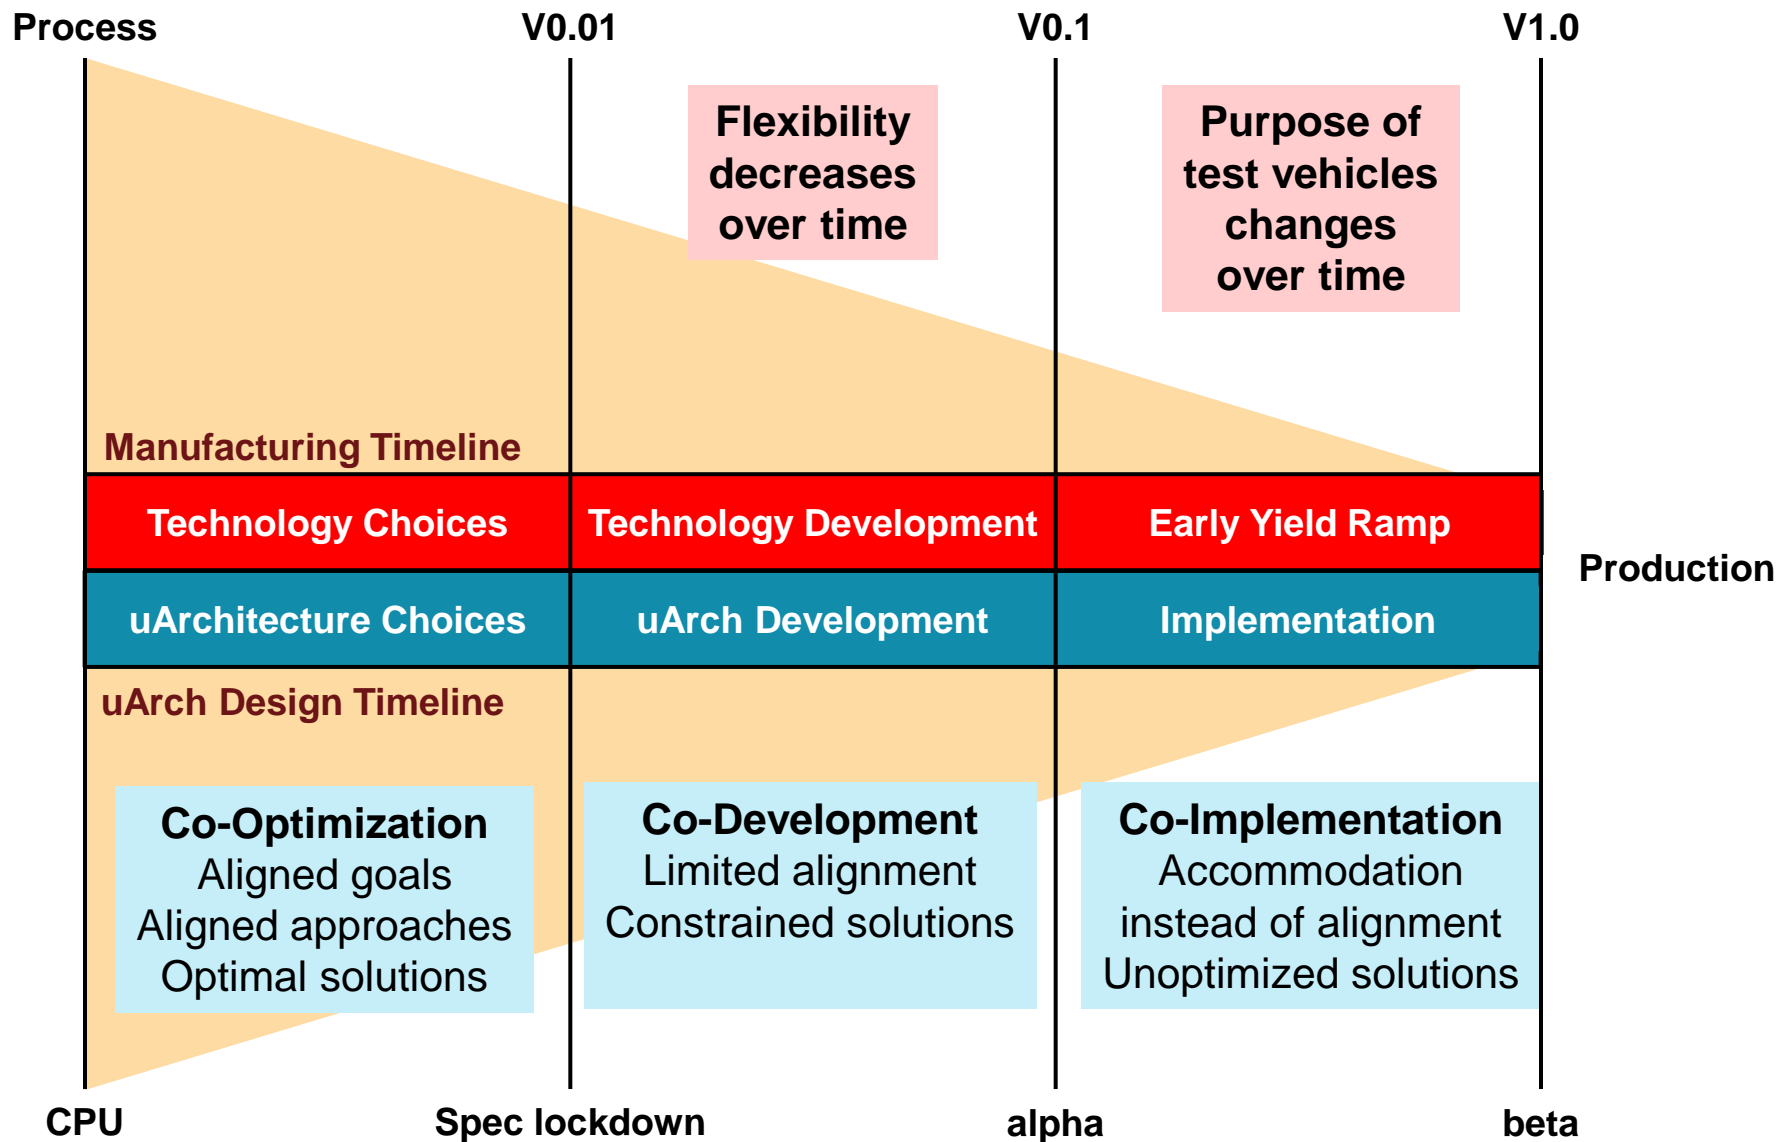
- 900+ cell options:
V_T, L, Drive Strength
- Reticle and Chip floorplanning
 - Placement
 - Routing
- Power delivery
- Statistical Timing
 - Random vs. Systematic
 - Correlation



Etch, CMP loading
Thermal loading
Power delivery network



Co-Optimization Timeline



ARM PDK – Development & Complexity

Predictive PDK

Electrical Models
(BSIM CMG)

DRC

LVS

Extraction
(ITF Standard)

TechLEF
(LEFDEF
Standard)

V_{th}

σ

Temperature
Dependence

Lithography
Dependence

FinFET
Parameter
Handling

Mask Layer
Derivations

Stack
Composition


R&C
Matching

Lithography
Dependence

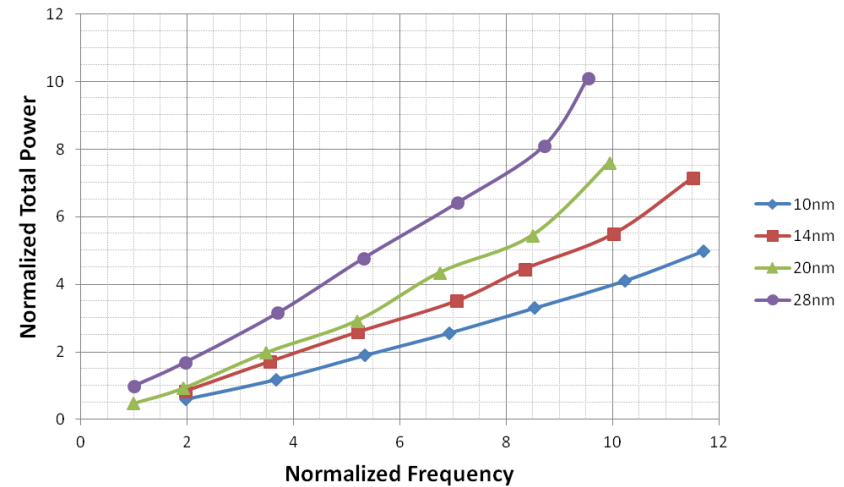
Predictive Library

- 2 basic litho options: SADP and LELELE
- 2 fin options: 3 fin and 4 fin
- Gives 4 combinations for library
- Can be used as-is to synthesize simple designs

Predictive Technology Benchmarking



Predictive PPA



Sources of Variability

- Lithography/Etch
 - Line edge roughness
 - CD variation
 - Influence of neighbors
- Device
 - Implant variation
 - Variation between N and P
 - Stress/strain effects
- Interconnect
 - Dielectric variation
 - Via/contact quality
 - Metal width/height variation
- Deterministic versus Random



Effects of Variability

$$I_d \approx \mu \cdot C_{ox} \left(\frac{W}{L} \right) (V_{gs} - V_t)^\alpha$$

- Leakage
 - Variation in L , V_t , μ , t_{ox}
- Performance
 - Changes in L , W , R , C , V_t , μ ,
- Min VDD
 - Changes in V_t , L , W
 - SRAM bit cell main limiter
- Dynamic power
 - Changes in C
 - Side effect of changes in performance, leakage
- Yield
 - Indirect result of others
 - Parameter goes beyond spec + tolerance




Margins

- Excess margins are wasteful
- Lack of margins is a bigger problem
- Leads to layers of margin, padding
 - Device
 - SPICE model
 - Standard cell
 - Static timing analysis
 - Manufacturing test
 - Etc. etc. etc.



Design Challenge: Margins



Design margins are barriers to efficiency

Variations:

Voltage fluctuates between 0.9V and 1.1V

PLL can jitter by 5ps

Silicon can be fast and leaky or slow

...

Uncertainties:

When exactly is my voltage 0.9V?

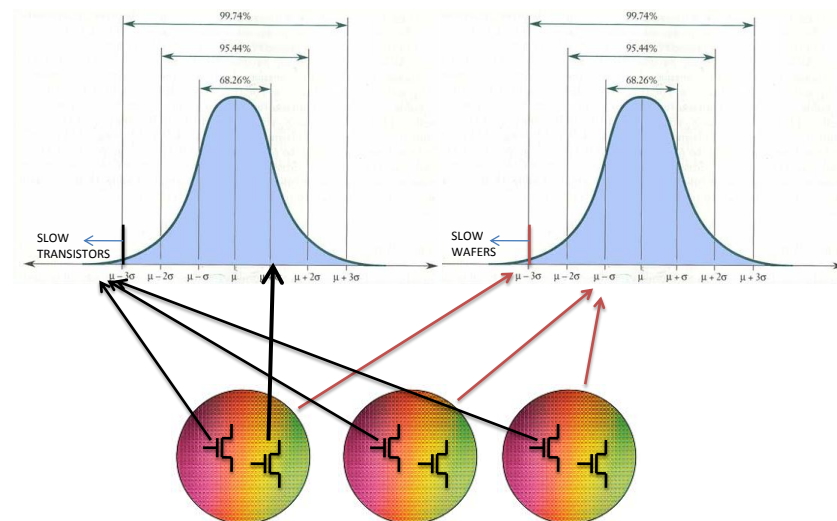
When exactly will I get jitter?

When will I get slow or fast silicon?

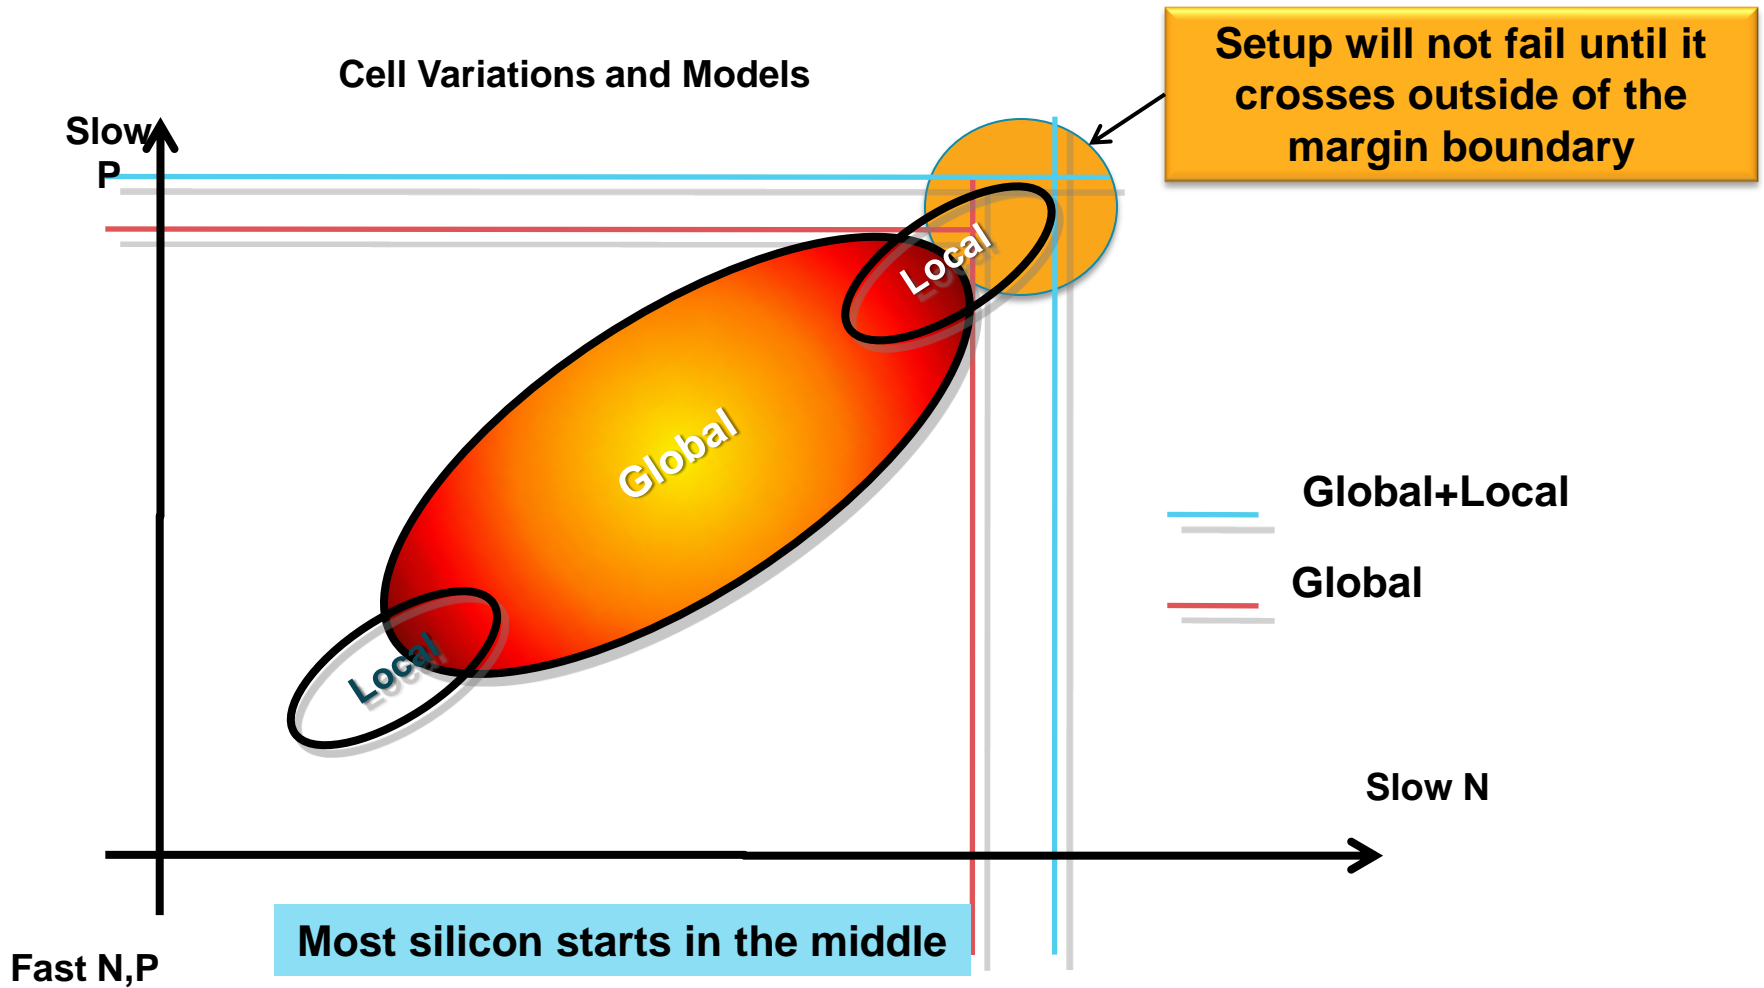
Uncertainty → margins
Excess margins are wasteful!

Process Corners and Margining

- “Everybody knows” that the SS and FF corners are 3 sigma
- Sigma of what?
 - Transistors across process
- Does this variation matter?
 - Yes, if you’re looking at WAT data.
 - No, if you care about chip behavior
- Need a better way to account for variation

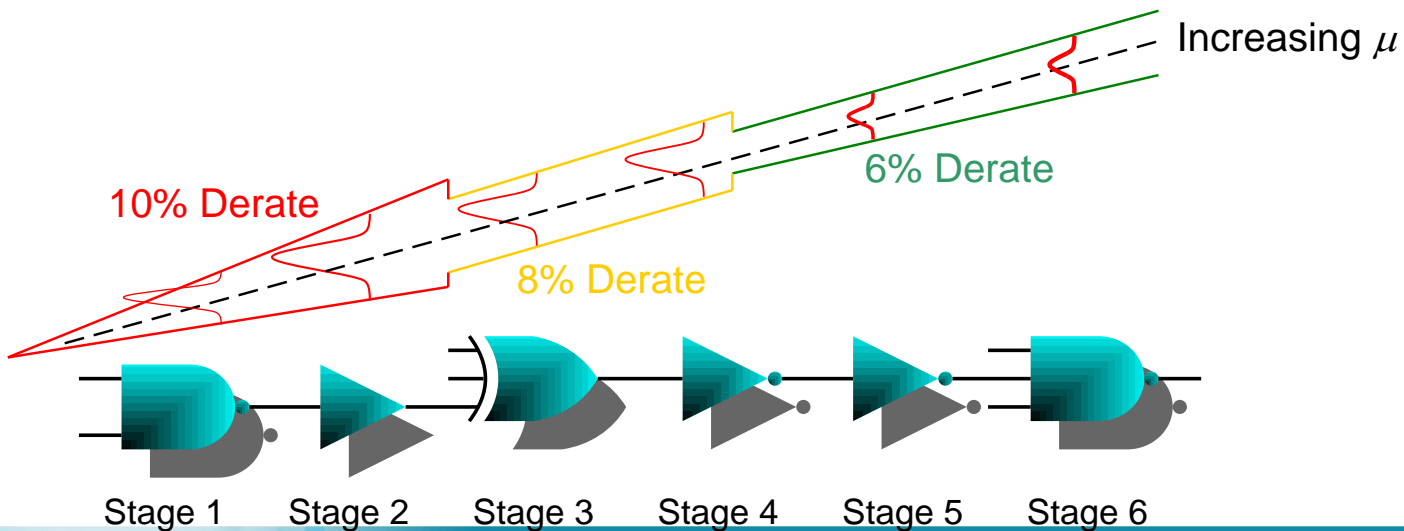
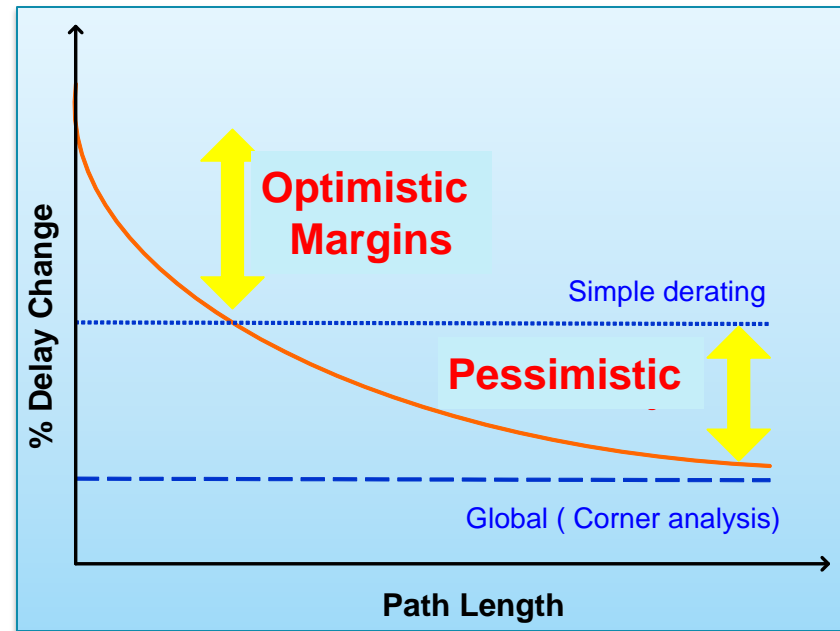


Delay Degradation ≠ Failure



Modeling On Chip Variation

- Random Variation
 - Derive OCV derate by Monte Carlo simulation
 - OCV derate \propto path variation / mean path delay
- Systematic Variation
 - Location based – derived using test-chips

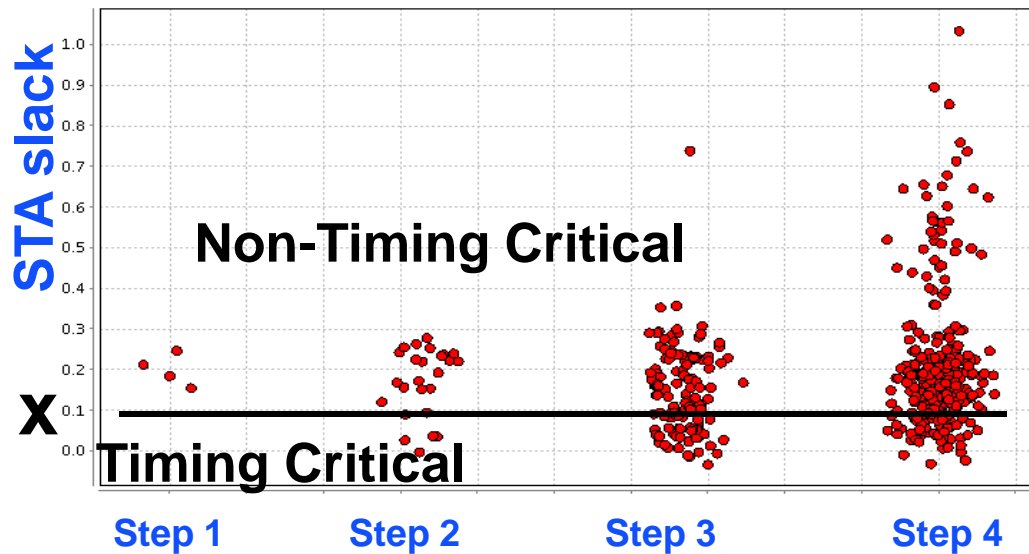


Finding Variation: Basic Test Axioms

- You can't find all the bad parts
- Every test will have unique fails
- Many parts with defects will work fine in the product
- Test environment is never the same as operating environment (temperature, voltage, clocking, noise)
- About 10% of any defect type are latent reliability problems
- Cost of a failure increases ~10X per step from wafer to package to board to system to product
- Better to reject a good part than ship a bad one
- Memory is complicated – use BIST
- Analog is complicated – measure subset of key parameters

At-Speed Scan Test and STA

- Key is to identify an appropriate WS test frequency for at-speed scan tests
 - Static timing analyzer not adequate at predicting true silicon timing
- STA not predicting order or paths



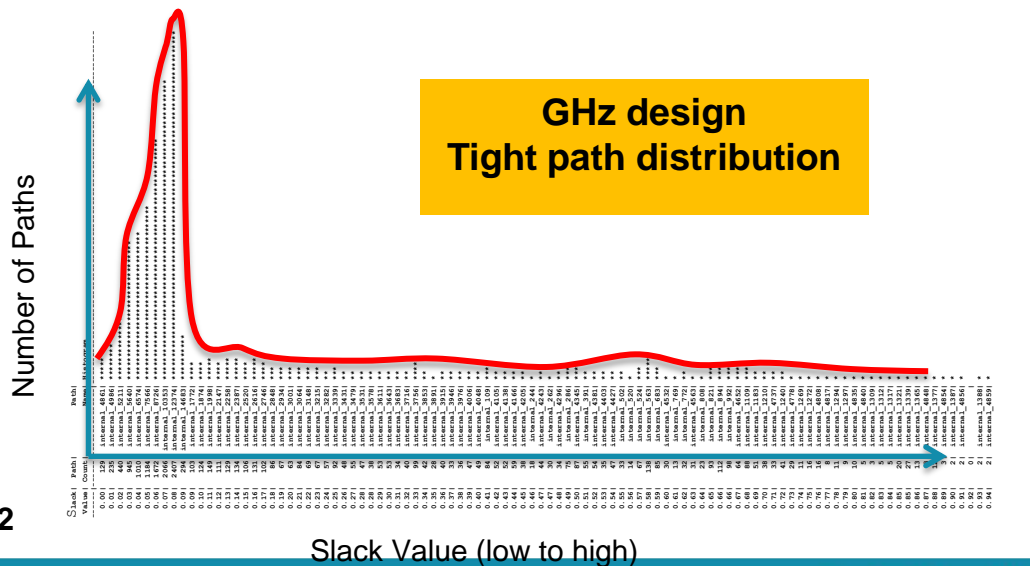
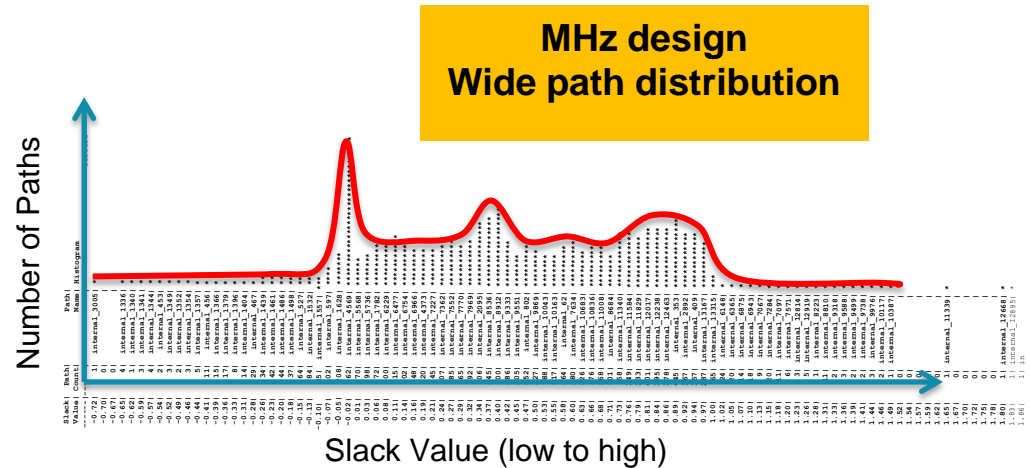
Janine Chen ,
Brendon Bolin,
Li-C Wang,
Jing Zeng,
Gagi Drmanac,
Michael Mateja,
“Mining AC Delay
Measurements for
Understanding
Speed-limiting
Paths,” ITC, 2010

M. Mateja, AMD, ITC 2011


Change in Frequency Distribution

- Higher speed design increases number of near critical paths
- These will fail in different order than STA predicts
- More complicated margining needed
- Increased susceptibility to local variation

T. McLaurin, IEEE D&T 2012

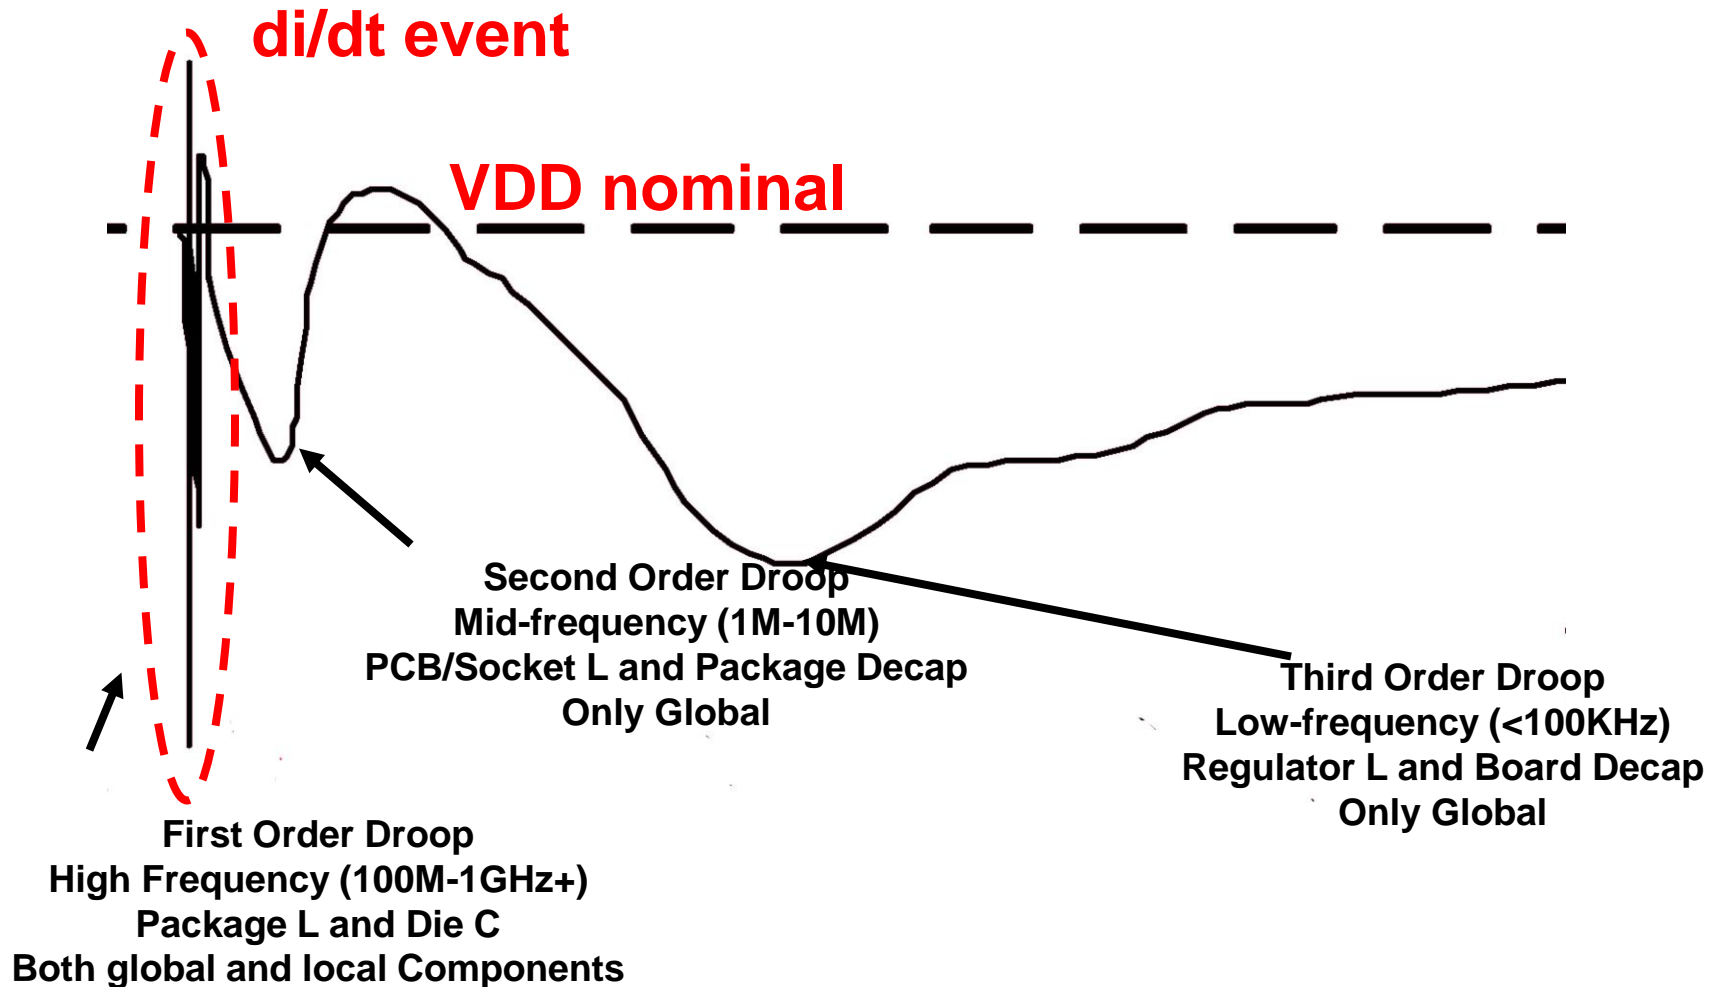


Challenge: Power Delivery




- For iso-power and VDD reduction for future process nodes current drawn will increase (V goes down, I goes up)
 - CV^2F constant \Rightarrow C or F can increase
 - More logic or run faster
- Delivering higher current will impact VDD integrity
 - IR due to increased steady-state current
 - Ldi/dt due to activity rate changes – function of ΔI and frequency
 - Timing impact higher at lower VDD

VDD Droops – Frequency Spectrum



On-die VDD Time-Domain Response to a $L di/dt$ emergency showing various frequency components
(Source: Muhtaroglu et al.)

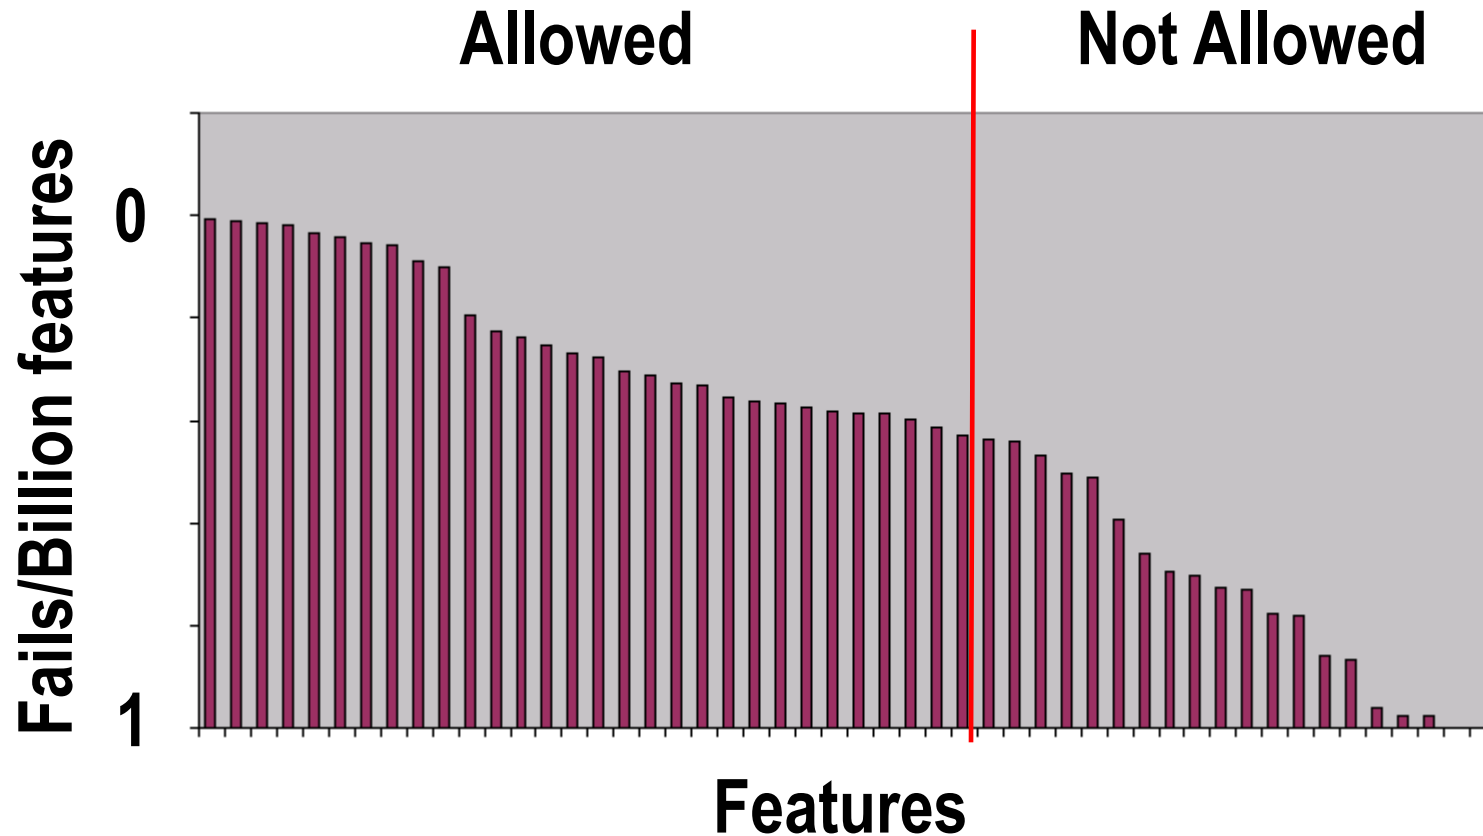
Challenge: Parametric Yield



- Assumes features allowed by the rules yield equally well.
- The design rules do not allow features that do not yield.

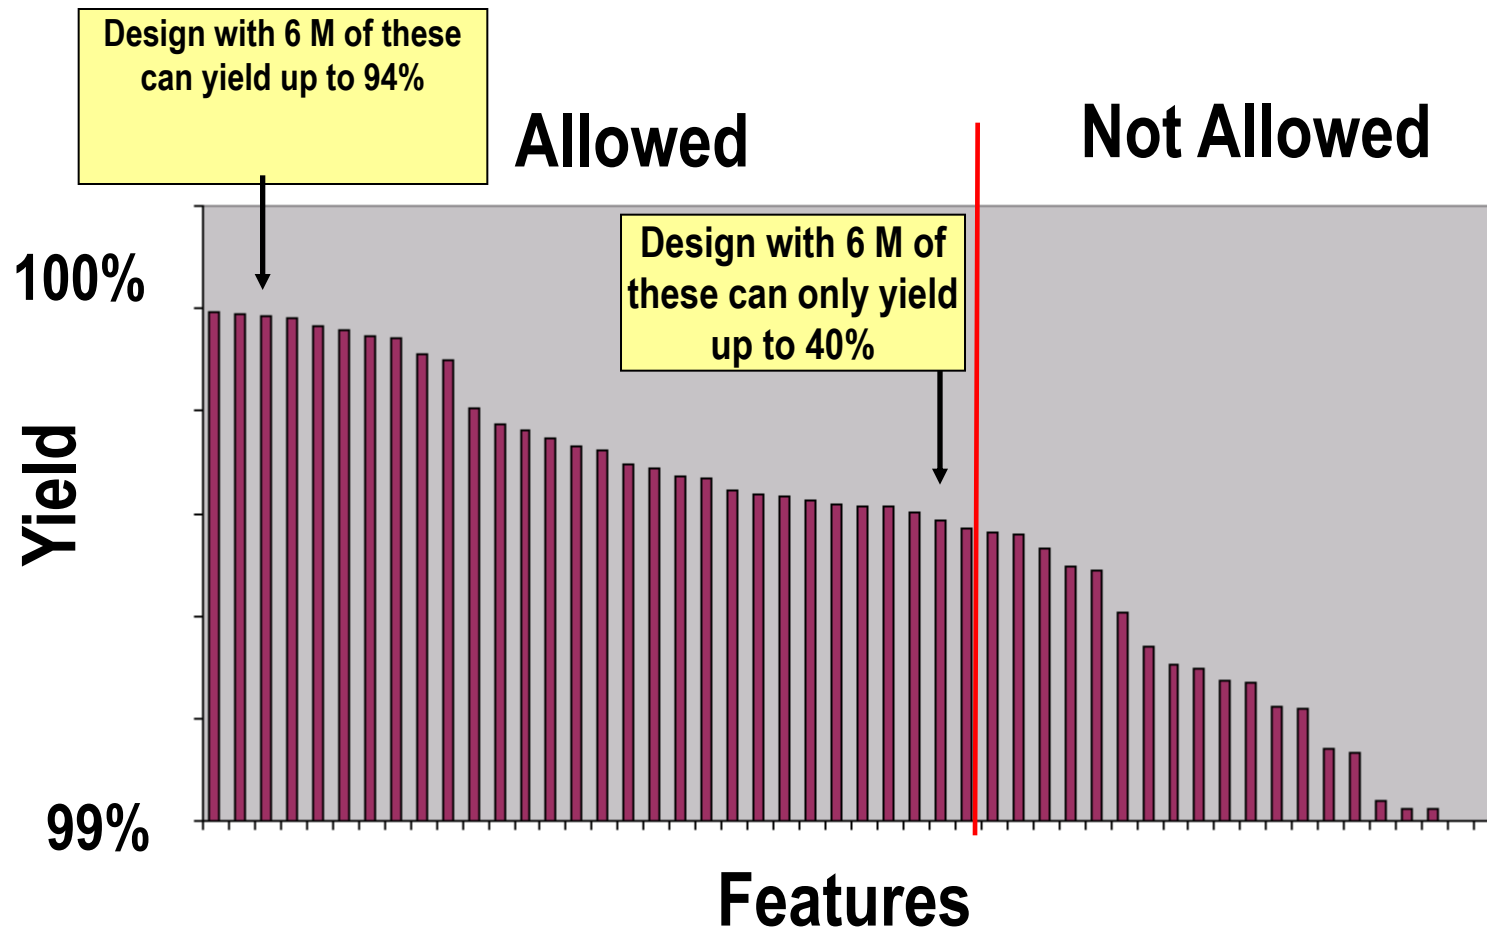
Source: Bob Madge

The True Feature Yield Curve



Source: Bob Madge

Feature Types Drive Product Yield



Source: Bob Madge

What's Ahead?

- The original questions remain
 1. How fast will a CPU go at process node X?
 2. How much power will it use?
 3. How big will it be?
- New ones are added
 4. How should cores be allocated?
 5. How should they communicate?
 6. What's the influence of software?
 7. What about X? (where X is EUV or DSA or III-V or Reliability or eNVM or IOT or)
- Collaboration across companies, disciplines, ecosystems

Fin

