A Roadmap for Low-Power Design: Trends, Technology, Tools

Andrew B. Kahng
UCSD CSE and ECE Departments

abk@ucsd.edu
http://vlsicad.ucsd.edu

Agenda

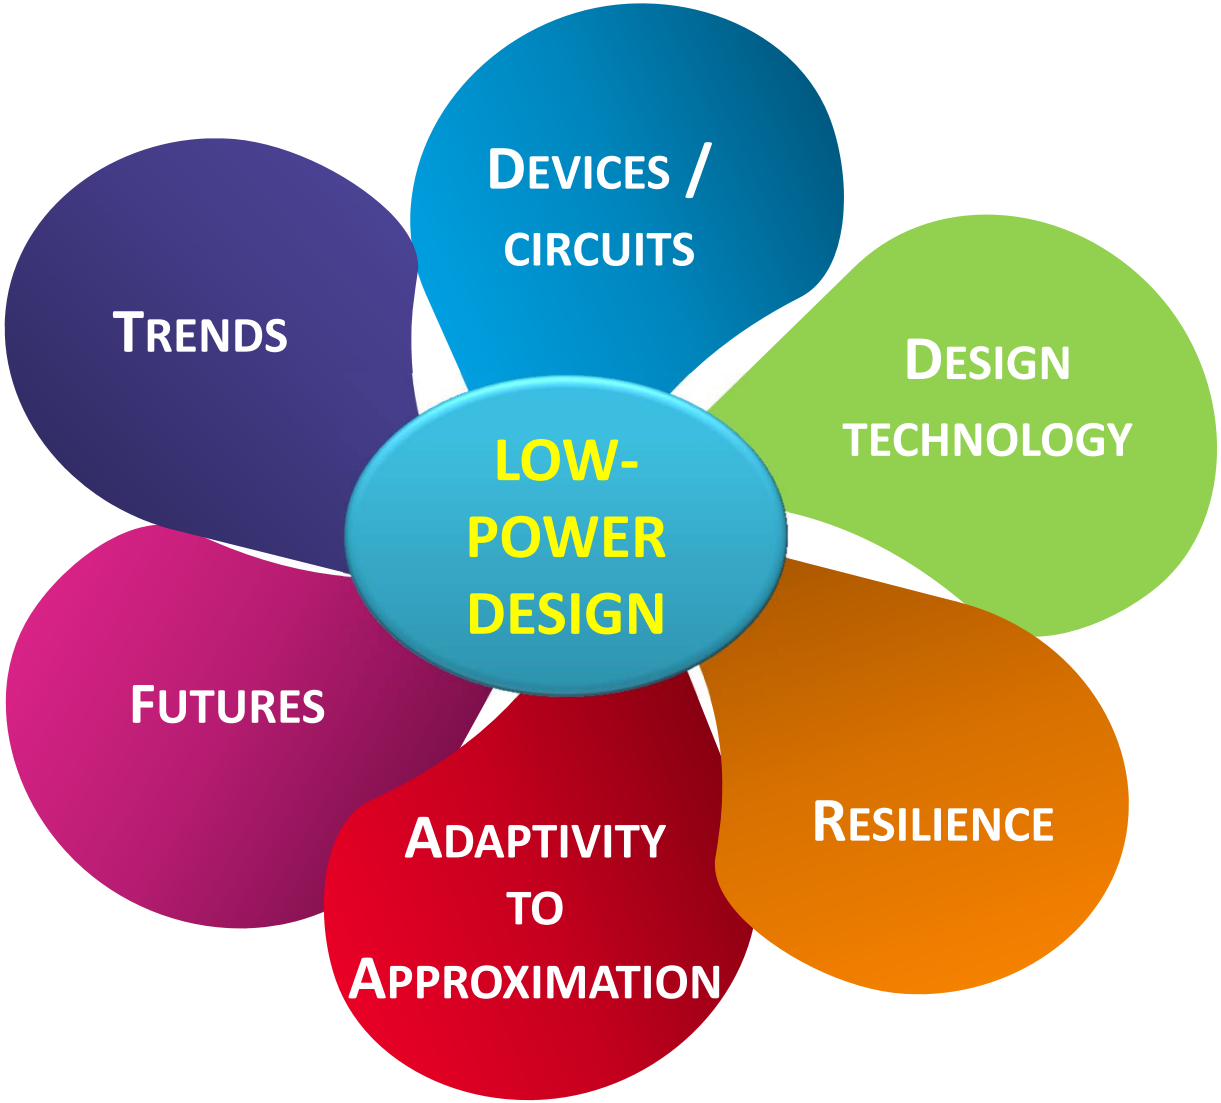
DEVICES / CIRCUITS
TRENDS
DESIGN TECHNOLOGY
LOW-POWER DESIGN
FUTURES
ADAPTIVITY TO APPROXIMATION
RESILIENCE

Agenda

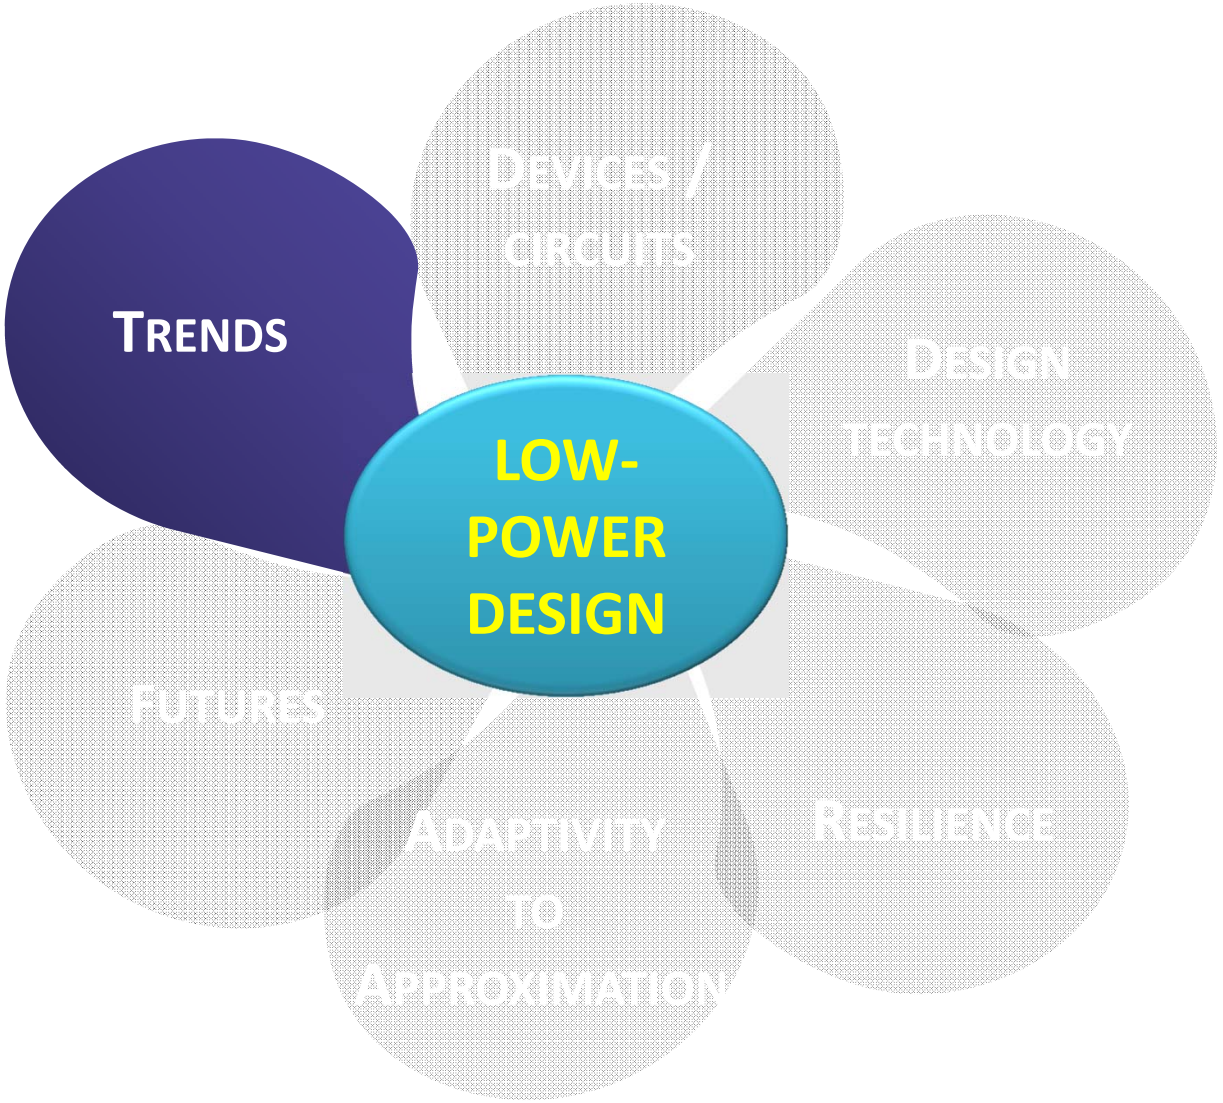
TRENDS
DEVICES / CIRCUITS
DESIGN TECHNOLOGY
LOW-POWER DESIGN
FUTURES
ADAPTIVITY TO APPROXIMATION
RESILIENCE

Trend 1: Race to the End of Roadmap

- $$$ for tech, design enablement: go big or go home
- Node pacing not slowing despite near-term “red bricks”
  - No EUV → multi-patterning
  - No Cu replacement → resistivity, rise of MOL / BEOL RC’s, variability
  - Reliability, layout restrictions → less benefit from node (20SOC “lost” like 45nm?)
  - Especially tough for early-adopter fabless
- Intrinsic mismatch of design-process time constants → margins!
  - Technology development, market definition, architectural design = O(years)
  - RTL-to-GDS implementation, reliability qualification = O(months)
  - Fab latency, cycles of yield learning, design re-spins = O(weeks)
  - Process tweaks, design ECOs = O(days)
  - Root cause of model-hardware miscorrelation, model guardbanding
- Paper to v1.0 SPICE models: 18 months → 12 months at N10
  - Will see how this works out…

Trend 2: Low Power Grand Challenge

- **Drivers for semi growth share critical requirement: LOW POWER**
  - Mobility
  - Big data, green datacenters, cloud
  - IoT
- Low-power design techniques increase design burden
- Added complexity of **system** + analysis + optimization
  - Multiple supply voltages
  - Multiple voltage domains
  - Extreme power, clock gating
  - DVFS
  - MTCMOS
  - Multi-Lgate
  - …

Power or Performance?

- **Cannot escape basic "shape" of tradeoff**
  - More power reduction (logic, Vt) available when freq ↓
  - ~Cubic relationship between power and frequency
- New designer mantras
  - "Highest performance at low power"
  - "Minimum V for any given throughput"

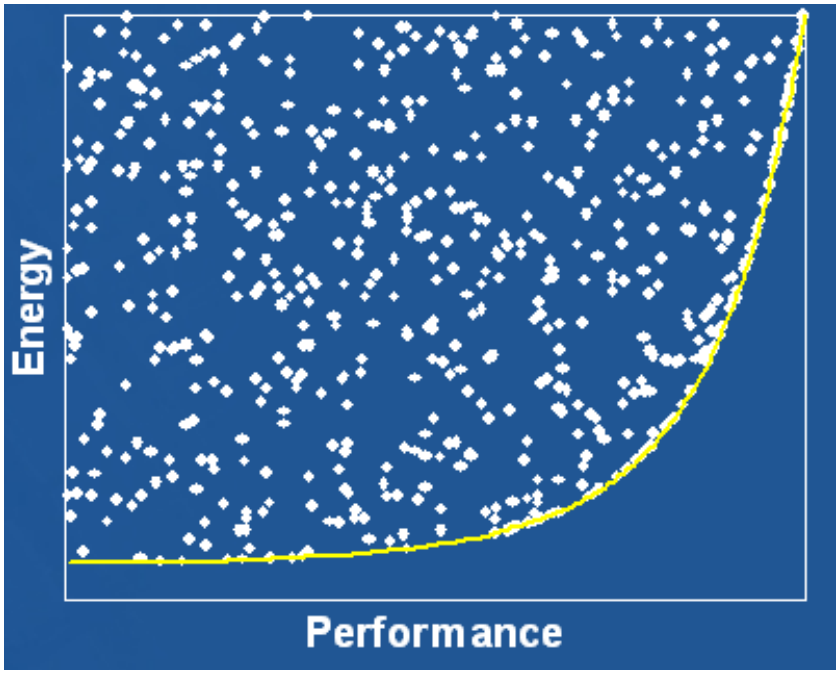

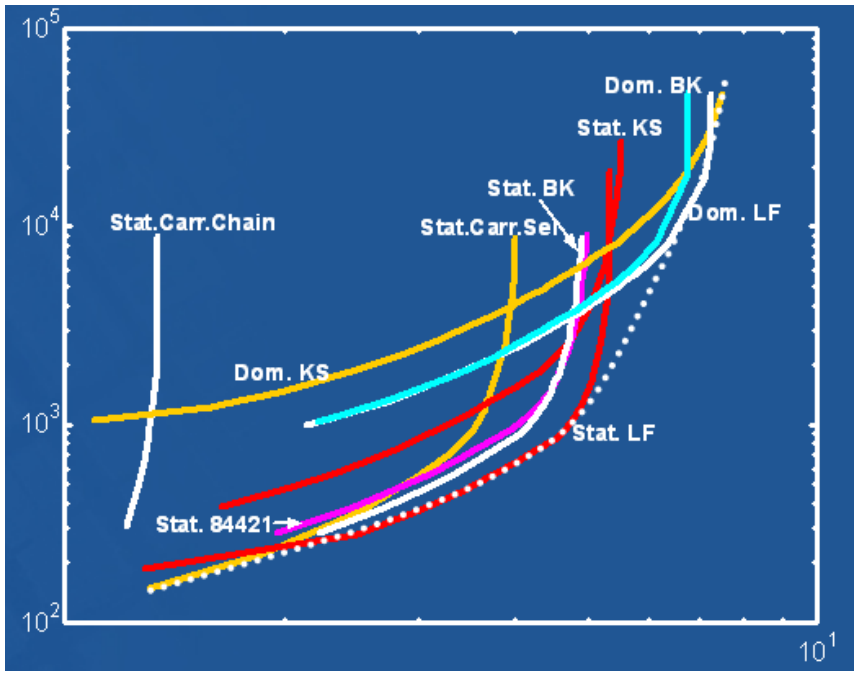

M. Horowitz, "Scaling, Power and the Future of CMOS", 2006 Workshop on On- and Off-Chip Interconnection Networks for Multicore Systems

Energy vs. Delay: Near-Threshold Computing?

- Supply voltage at near-threshold region

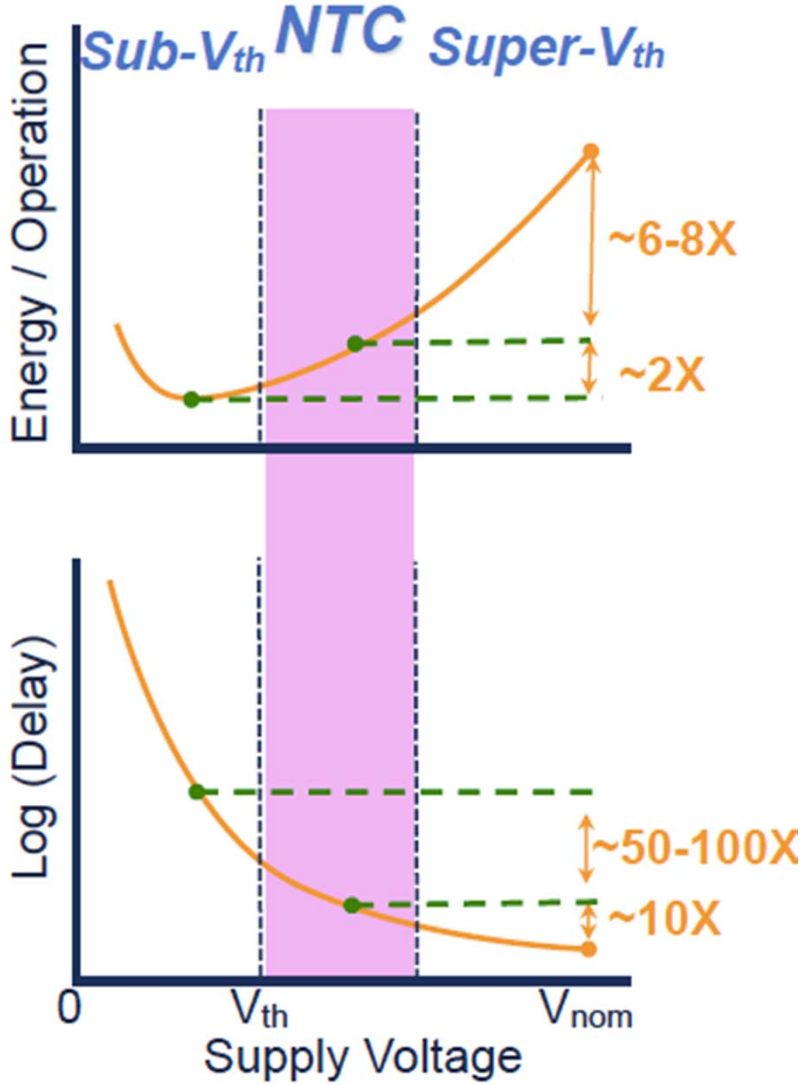

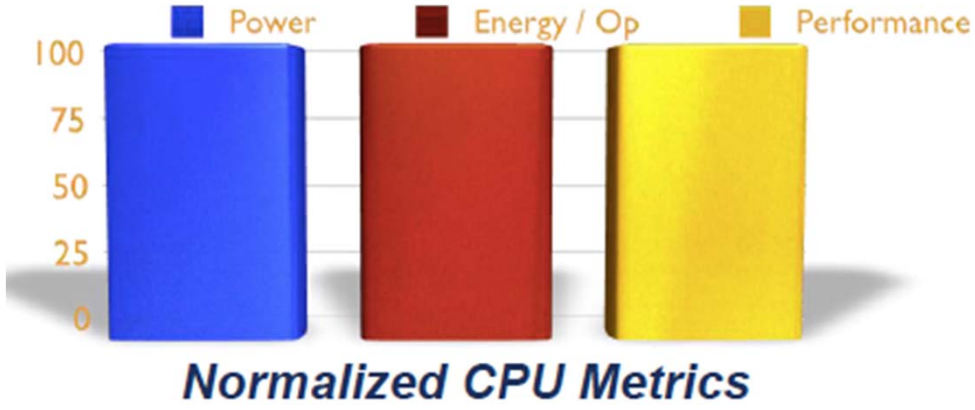

>60X power reduction
6-8X energy reduction
Enables 3D integration

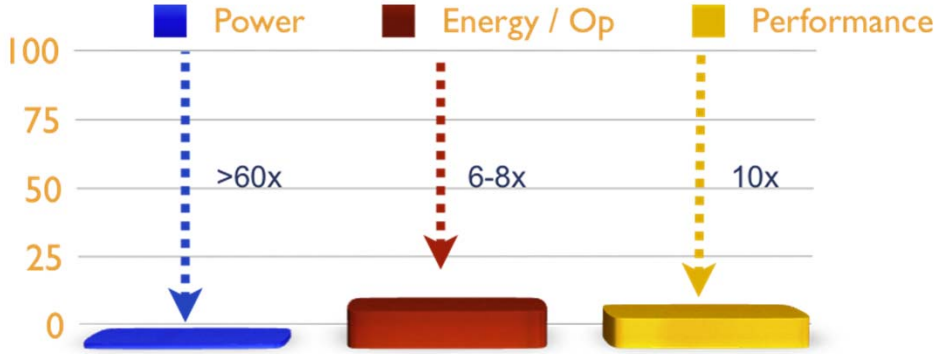

R. G. Dreslinski, et al., "Centip3De: A 64-Core, 3D Stacked Near-Threshold System", HotChips 2012.

NTC Has Barriers …

- Performance loss
  - 45nm FO4 delay at NTC supply of 400mV vs. 1.1V: 10X slower
- Increased performance variation
  - Performance variation due to global process variation alone: ~30% at 1.1V, up to ~400% at 400mV
- Increased functional failure
  - Random dopant fluctuation, line edge roughness → positive feedback of device mismatch in SRAM

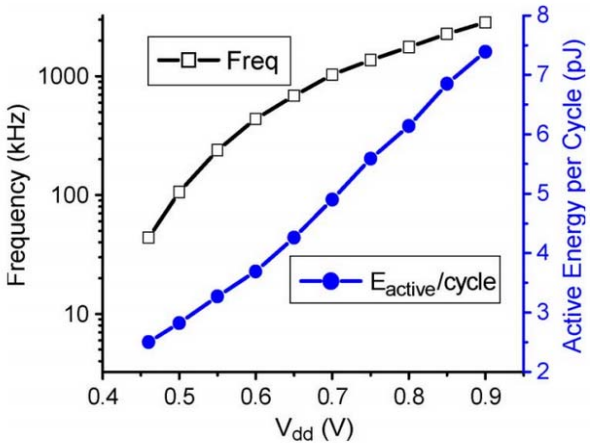

Phoenix frequency and energy breakdowns at various supply voltages.

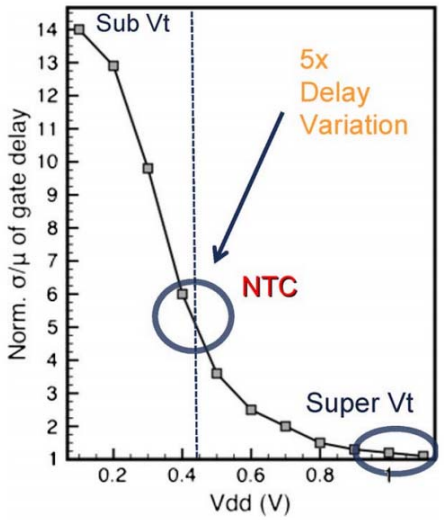

Impact of voltage scaling on gate delay variation.

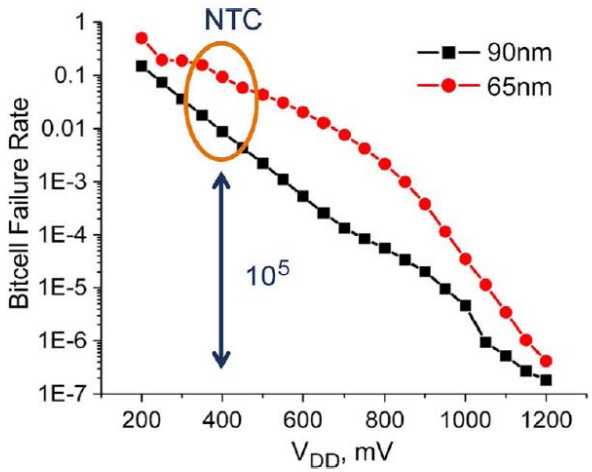

Impact of voltage scaling on SRAM failure rates.

R. G. Dreslinski, et al., "Near-Threshold Computing: Reclaiming Moore's Law Through Energy Efficient Integrated Circuits", Proceedings of the IEEE 98(20) (2010). Pp. 253-266.

… Which Have Workarounds, but

- Performance loss
  - Cluster-based architecture
  - Device optimization
- Performance Variation
  - Soft-edge clocking
  - Body biasing
- Functional Failure
  - Alternative SRAM Cells
  - SRAM robustness analysis
  - Reconfigurable cache designs
- ***(FinFET changes the context!)***

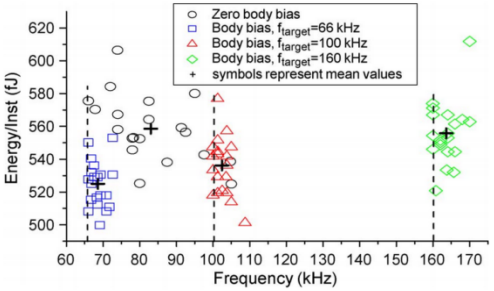

Body Biasing tech for three target frequencies

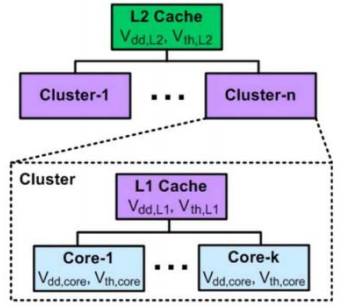

Cluster-based architecturevb

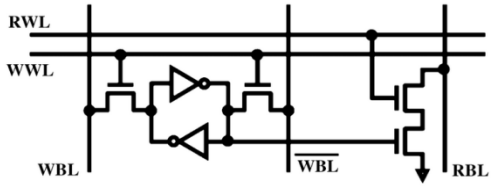

Alternative 8 T SRAM cell, decoupling the read and write.

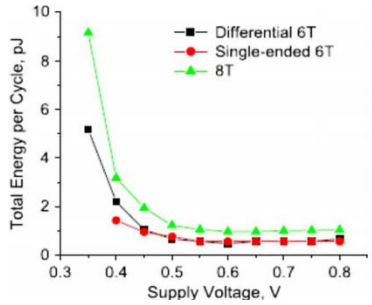

Energy of SRAM topologies for 20-cycle L2 cache across voltages

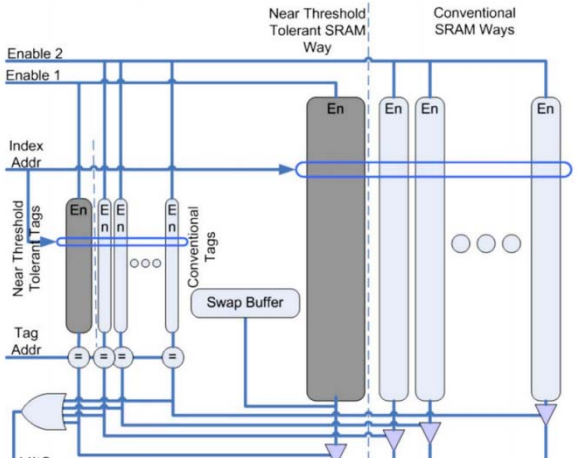

Alternative L1 cache design with one cache way NTC enabled

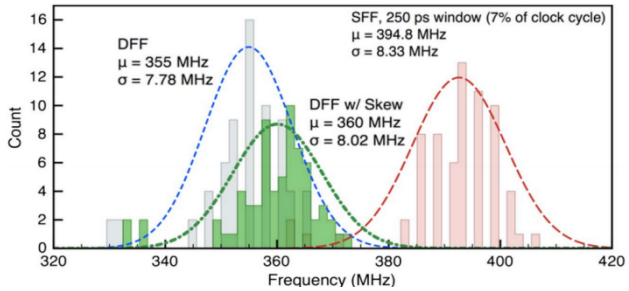

FIR filter with soft edge clocking compared to standard flip-flops (SFF); presented with and without useful skew.

R. G. Dreslinski, et al., "Near-Threshold Computing: Reclaiming Moore's Law Through Energy Efficient Integrated Circuits", Proceedings of the IEEE 98(20) (2010). Pp. 253-266.

Agenda

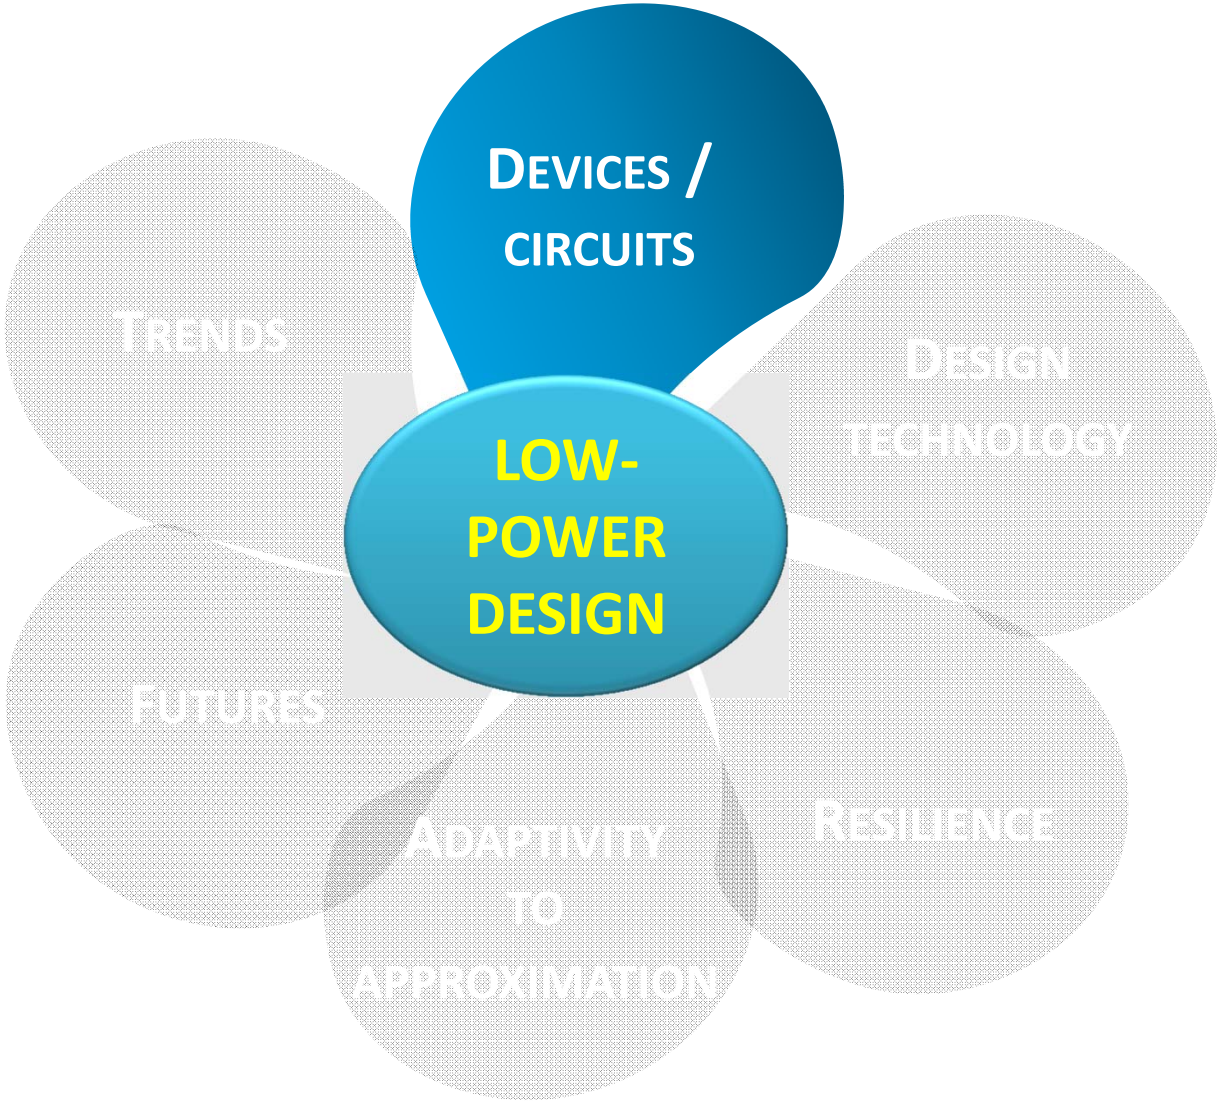
DEVICES / CIRCUITS
TRENDS
DESIGN TECHNOLOGY
LOW-POWER DESIGN
FUTURES
RESILIENCE
ADAPTIVITY TO APPROXIMATION

Trend 3: High-Value Equivalent Scaling

- Device roadmap (FinFET, FDSOI) helps enormously
  - Better electrostatic control
  - Lower leakage current

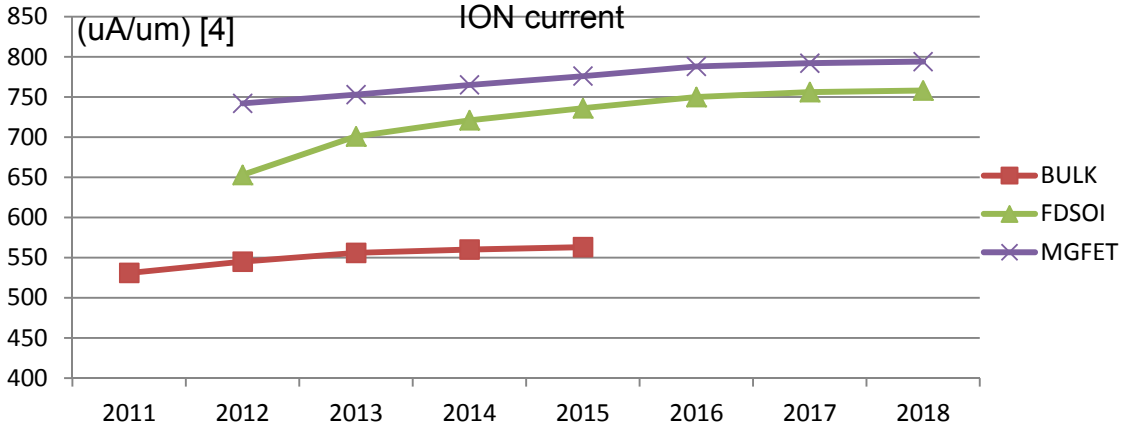

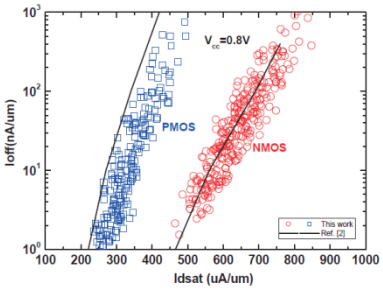

TSMC 28nm planar bulk [1]

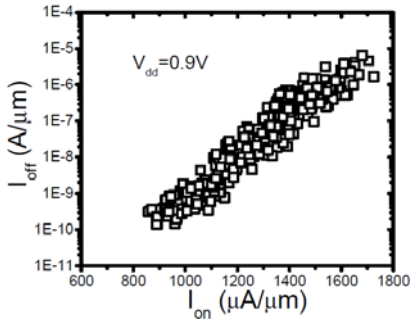

ST 20nm FDSOI [2]

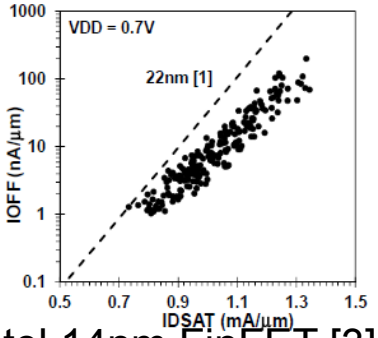

Intel 14nm FinFET [3]

[1] http://www.realworldtech.com/iedm-2010/6/
[2] S. Natarajan, et al., IEDM 2014, pp. 71-74.
[3] Q. Liu, et al., IEDM 2014. pp. 219-222.
[4] http://www.itrs.net

FinFET vs. FDSOI for Low-Power Design

	FinFET	FDSOI
Surface passivation	Yes	Yes
HK metal gate stack	Yes	Yes
BG biasing	No	Yes
DIBL	70mV/V	80mV/V

Multi-Vt techniques

- FinFET: better subthreshold swing, DIBL [Yeh 10]
  - → Performance less sensitive to Vdd
  - → lower Vdd, less active power at same speed
- FDSOI: more V_t control options [Skotnicki10] [PachaASX06][Biesemans]
  - → Metal gate stack changes work function, V_t
  - → Back-plane/gate doping
  - → Back-gate biasing

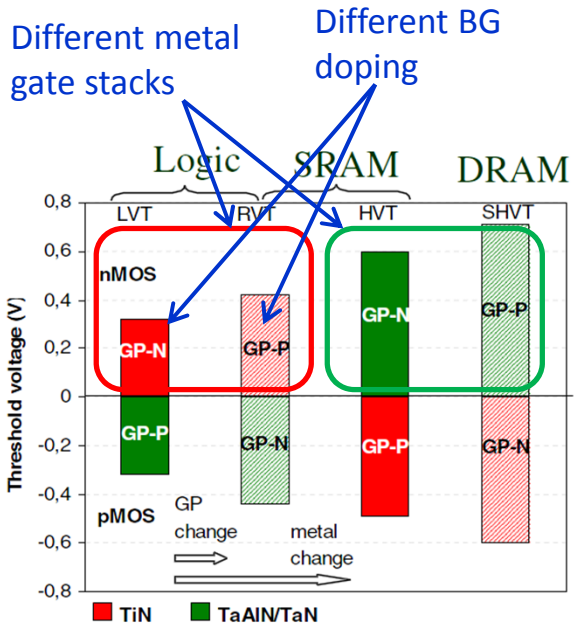

Much Wider Operating Voltage Range

- Supply voltage scaling is key low-power technique
- Enabled with FinFET, UTBB FDSOI
- Complexity explosion: modes, corners in timing closure

Technology	28nm UTBB FDSOI	28nm LP Bulk	32nm Bulk	22nm Trigate	**28nm UTBB FDSOI**
VDD operating range	0.6V-1.2V	0.34V-1V	0.28V-1.2V	0.28V-1.1V	**0.39V-1.3V**
Max measured Frequency	2.6GHz@1.3V	587MHz@1V	915MHz@1,2V	2.5GHz@1,1V	**2.6GHz@1.3V**
Frequency @Min voltage	1GHz@0.6V	3.6MHz@0,4V	3MHz@0.28V	16.8MHz@0.28V	**460MHz@0.397V**
Total power consumption	na	113mW@1V	400mW@1V	227mW@1V	**370mW@1V**
Peak energy efficiency	na	na	170pJ/cycle @0.45V	585GOPS/W @260mV	**62pJ/op @0.53V**

R. Wilson, et al., " A 460MHz at 397mV, 2.6GHz at 1.3V, 32b VLIW DSP, embedding FMAX tracking", ISSCC 2014, pp. 452-454.

Trends 1, 2, 3 → Design Closure Nightmare

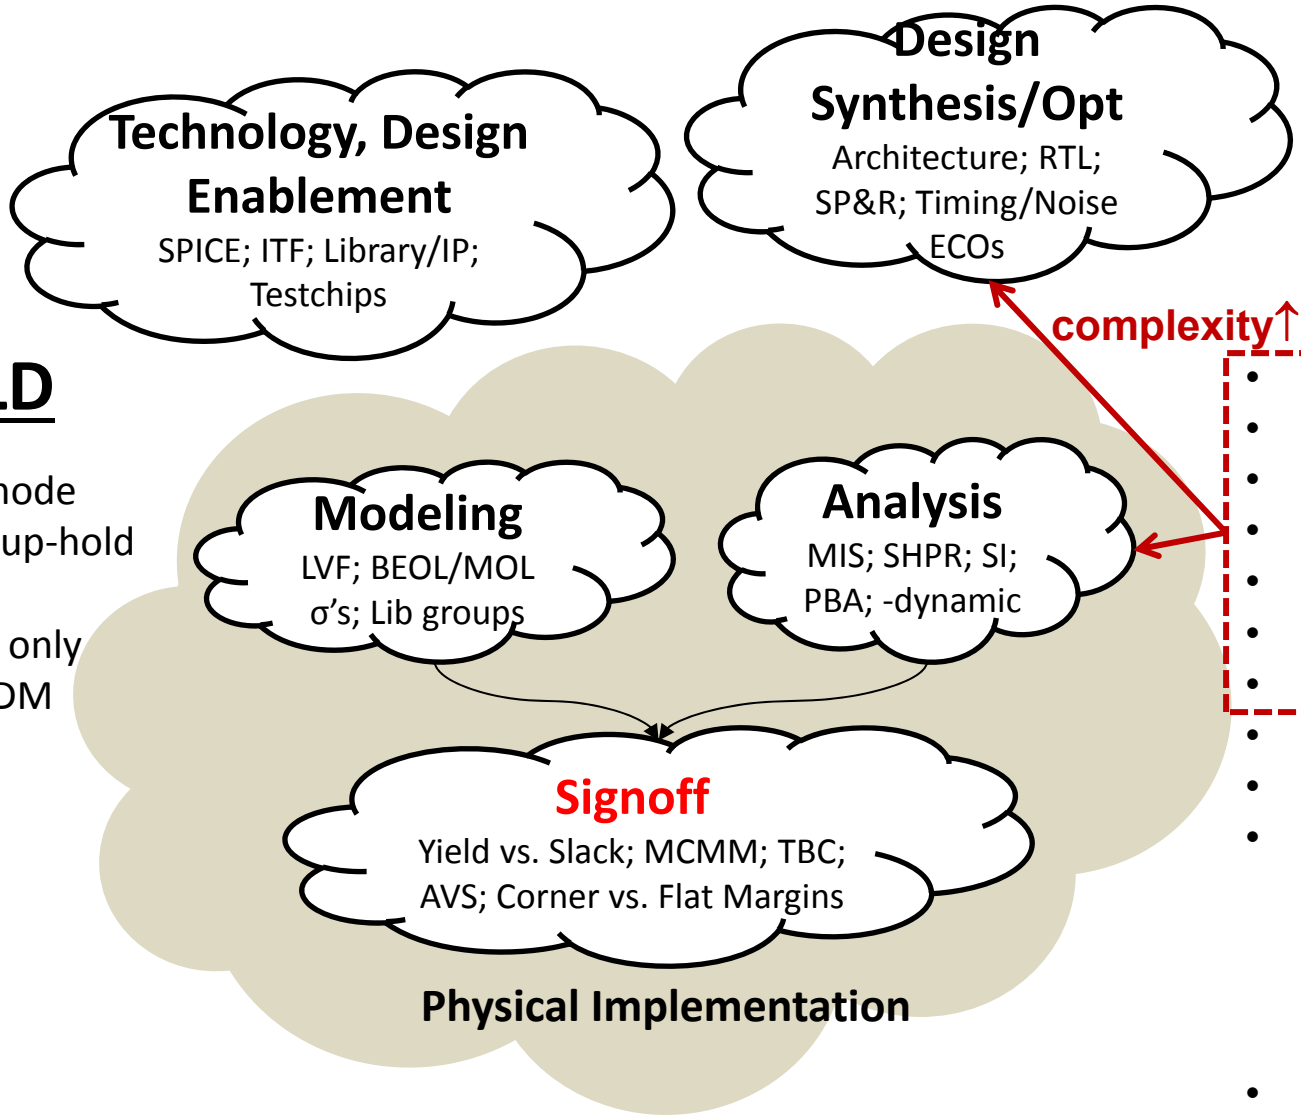

OLD

- 1 mode
- Setup-hold
- SI
- Cw only
- NLDM

NEW

complexity↑

- MCMM
- AVS
- Power & Clock gating
- Multiple voltage domains
- DVFS
- MTCMOS
- Multi-Lgate
- Cell-POCV / LVF
- Dynamic IR
- Wide/exploding corners, corner reduction, cross-corners (BEOL Cw, Ccw, RCw, temp, VDD)
- Flat margin selection
- Noise closure
- Aging/AVS

Agenda

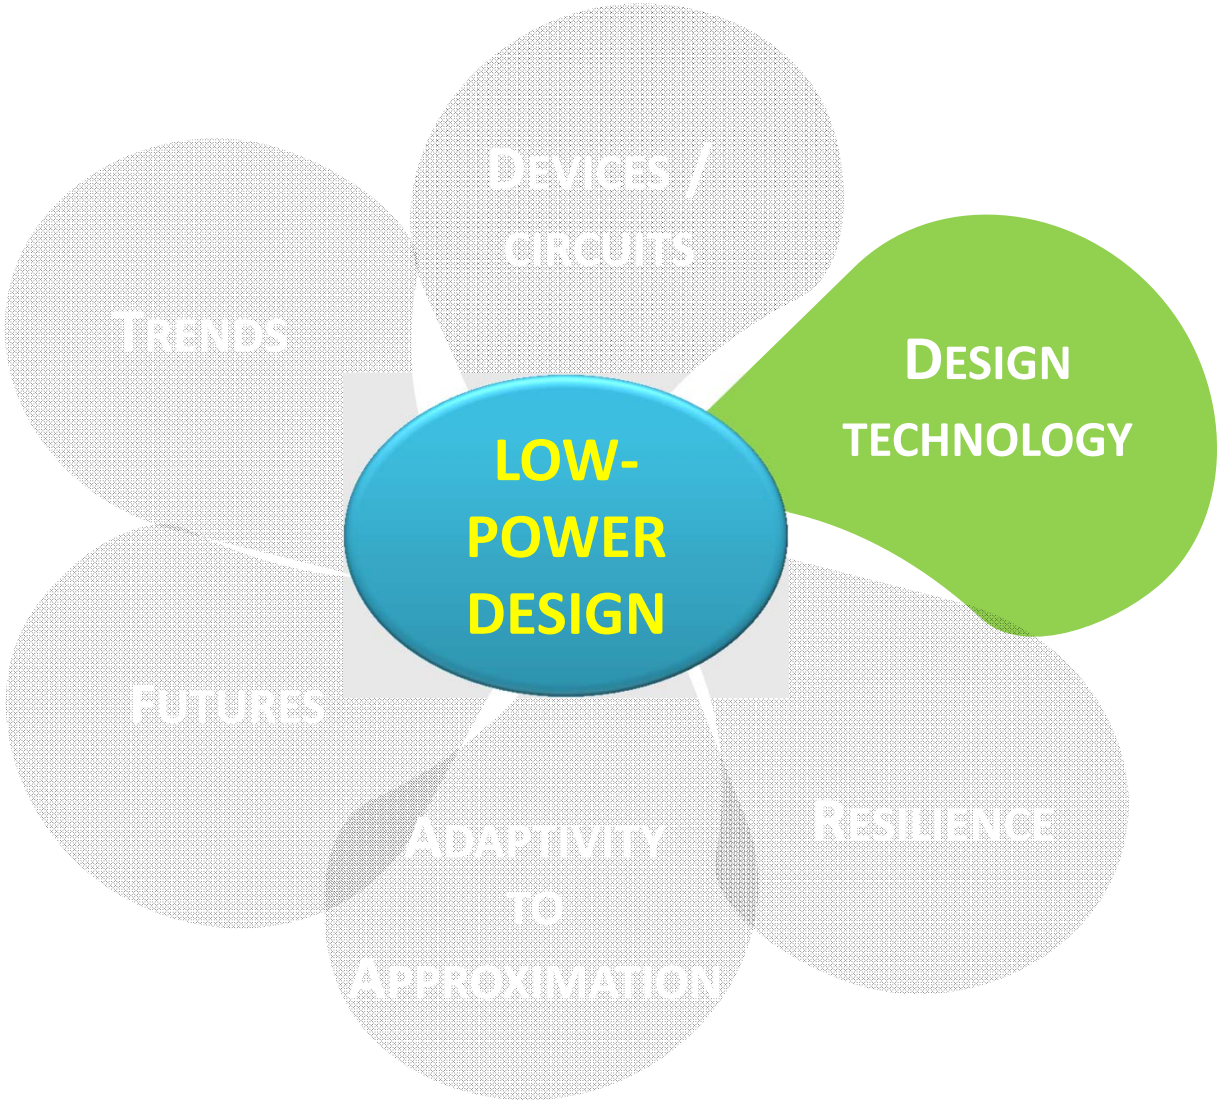
DEVICES / CIRCUITS
TRENDS
DESIGN TECHNOLOGY
LOW-POWER DESIGN
FUTURES
ADAPTIVITY TO APPROXIMATION
RESILIENCE

Alphabet Soup: DVFS, DCVS, AVS, PVS, SVS…

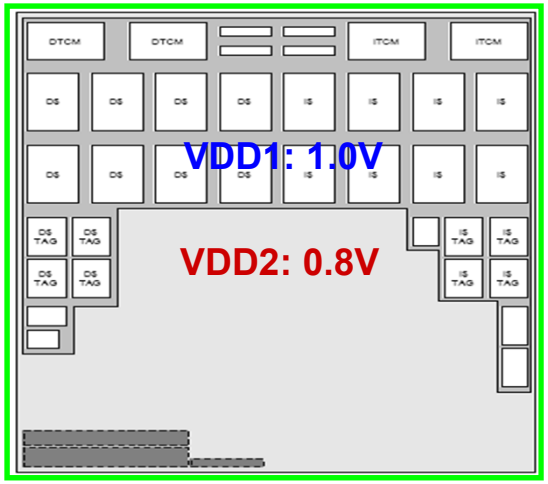

Multiple Voltage Islands (Multi-VDD)

Operating voltage is different, but, voltage has the fixed value

Dynamic Voltage / Frequency Scaling

Operating voltage and frequency change according to workload

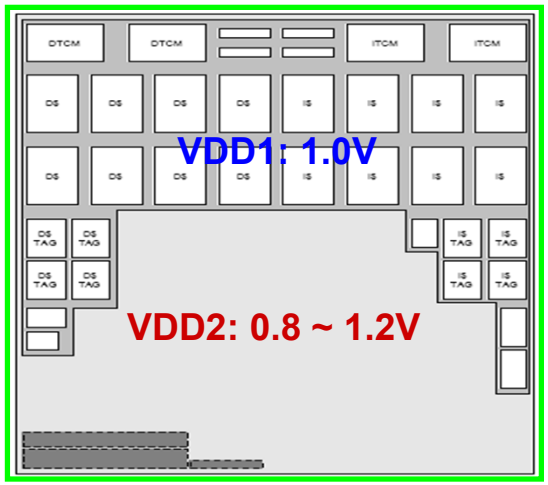

Adaptive Voltage Scaling

Operating voltage can vary to meet the performance requirement

(also used to compensate process variation)

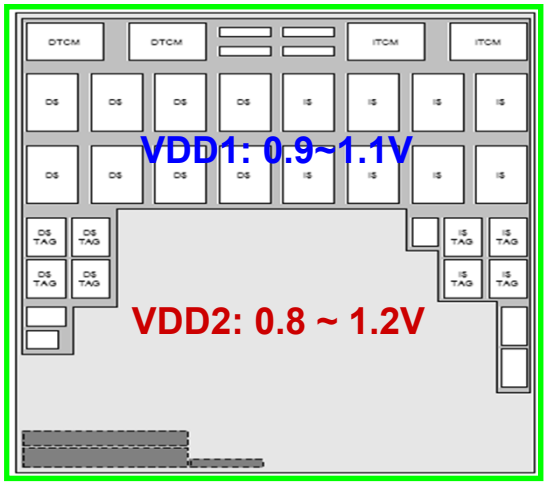

Low-Power Design Roadmap in the ITRS

ITRS: power and energy = the grand challenge for semiconductor industry

→ low-power design roadmap (2011)

Design Technology Improvement	*Year*	*Improvements*		*Description*
		Dynamic	*Static*	
Low Power Physical Libraries	Before 2011	1.50	1.50	Optimizing transistor size, layout style and cell topology for the standard-cell library
Back Biasing		1.00	1.35	Biasing wells of devices independently of the sources to shift the threshold voltage
Adaptive Body Biasing (ABB)		1.20	2.00	Delivering a positive or negative voltage below a transistor to reduce leakage
Power Gating		0.90	10.00	Turning off the power supplies to idle blocks for leakage reduction
Dynamic Voltage/Frequency Scaling (DVFS)		1.50	1.00	Dynamic management of supply voltage and operating frequency for power reduction
Multilevel Cache Architecture		1.00	1.20	Reduce amount of off-chip memory accesses for performance improvement and power reduction
Hardware Multithreading		1.00	1.30	Using multithreads to improve hardware utilization with leakage reduction
Hardware Virtualization		1.00	1.20	Using one physical server to support multiple guest operating systems simultaneously
Superscalar Architecture		1.00	2.00	Parallel instruction issueing and executing for performance improvement and power reduction
Symmetric Multiple Processing (SMP)		1.50	1.00	Lowering the frequency by using multiple processors and the parallel programming
Software Virtual Prototype	2011	1.23	1.20	Allow the programmer to develop software prior to silicon
Frequency Islands	2013	1.26	1.00	Designing blocks that operate at different frequencies
Near-Threshold Computing	2015	1.23	0.80	Lowering Vdd to 400 - 500 mV
Hardware/Software Co-Partitioning	2017	1.18	1.00	Hardware/software partitioning at the behavioral level based on power
Heterogeneous Parallel Processing (AMP)	2019	1.18	1.00	Using multiple types of processors in a parallel computing architecture
Many Core Software Development Tools	2021	1.20	1.00	Using multiple types of processors in a parallel computing architecture
Power-Aware Software	2023	1.21	1.00	Developing software using power consumption as a parameter
Asynchronous Design	2025	1.21	1.00	Total Non-clock driven design
Total		4.66	0.96	

Value of Low-Power Design Technology

Design Technology Improvement	*Year*	*Improvements*		*Description*
		Dynamic	*Static*	
Software Virtual Prototype	**2011**	**1.23**	**1.20**	Allow the programmer to develop software prior to silicon
Frequency Islands	**2013**	**1.26**	**1.00**	De fre
Extreme Power Gating	**2015**	**0.90**	**1.00**	Sh
Hardware/Software Co-Partitioning	**2017**	**1.18**	**1.00**	Ha
Heterogeneous Parallel Processing (AMP)	**2019**	**1.18**	**1.00**	pa
Many Core Software Development Tools	**2021**	**1.20**	**1.00**	Us pa
Power-Aware Software	**2023**	**1.21**	**1.00**	De co
Near-Threshold Computing	**2025**	**1.23**	**0.80**	Lowering Vdd to 400 - 500 mV
Asynchronous Design	**2027**	**1.21**	**1.00**	Total Non-clock driven design
Total		**3.47**	**0.96**	

SOC consumer portable chip in 2028:

- **2.5 billion logic gates**
- **Low-power DT reduces power from 48.8W to 9.1W**

Roadmap: Big Gaps Ahead

- ITRS: power of mobile SOC-CP driver keeps increasing…
- … even if envisioned low-power innovations are developed and deployed on time

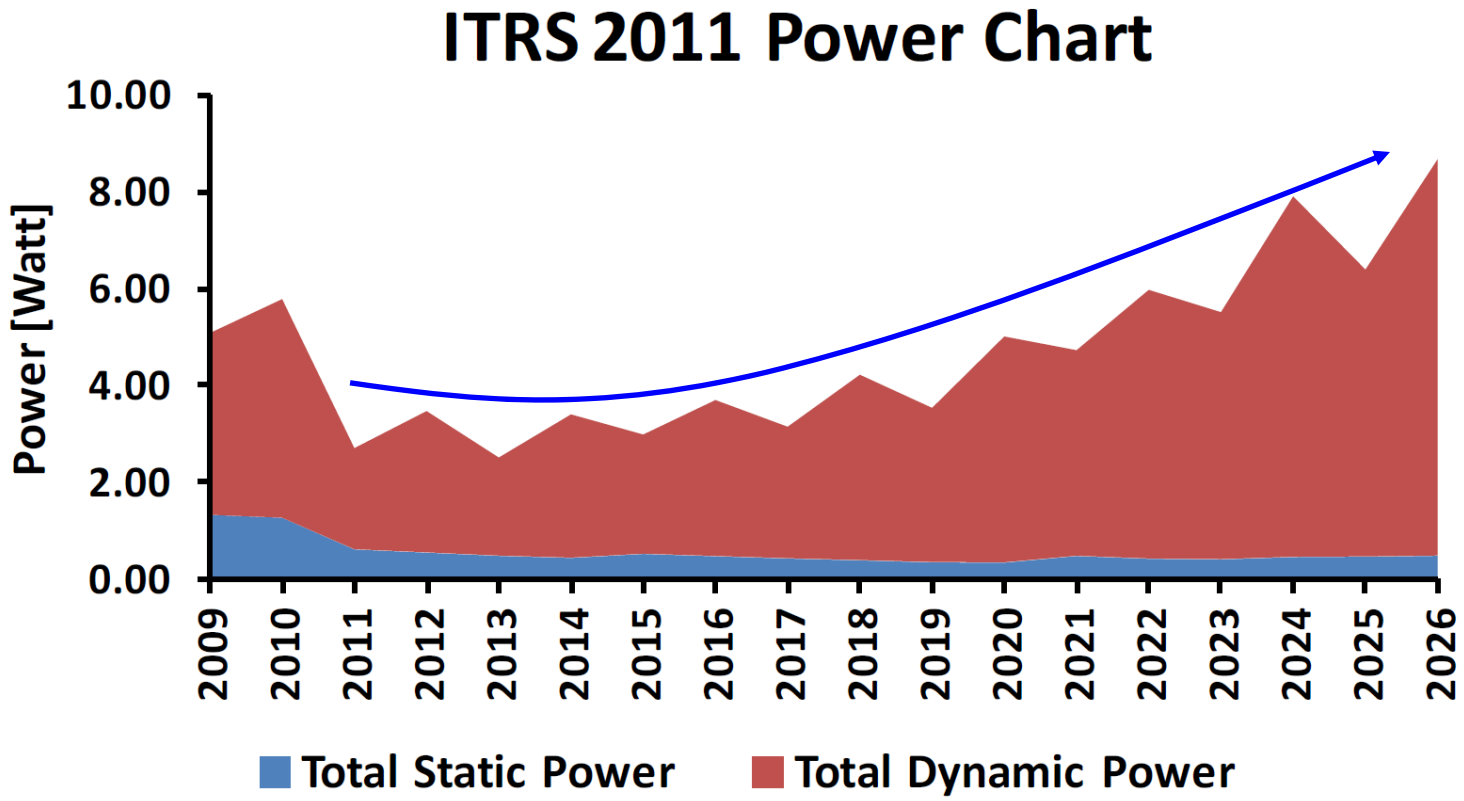

ADDED in 2013 Low-Power DT Roadmap

Approximate Computing

- Variable-accuracy computing (e.g., flexibly from 64b ↔16b)

4D Computing

- Reconfiguration on the fly

Adaptivity

- Recapture overdesign from wearout, variation margins

Power Gating Replacement

- HVT device as power switch hits headroom, area wall → ?

Extreme Heterogeneity

- "coprocessor-dominated architectures"
  - (pervasive heterogeneity; energy-efficiency from specialization; HW accelerators)

Extreme Power Gating

- Reaching the limits of shut-off

Signoff At Typical

- Use adaptivity to recover margin, overdesign

Agenda

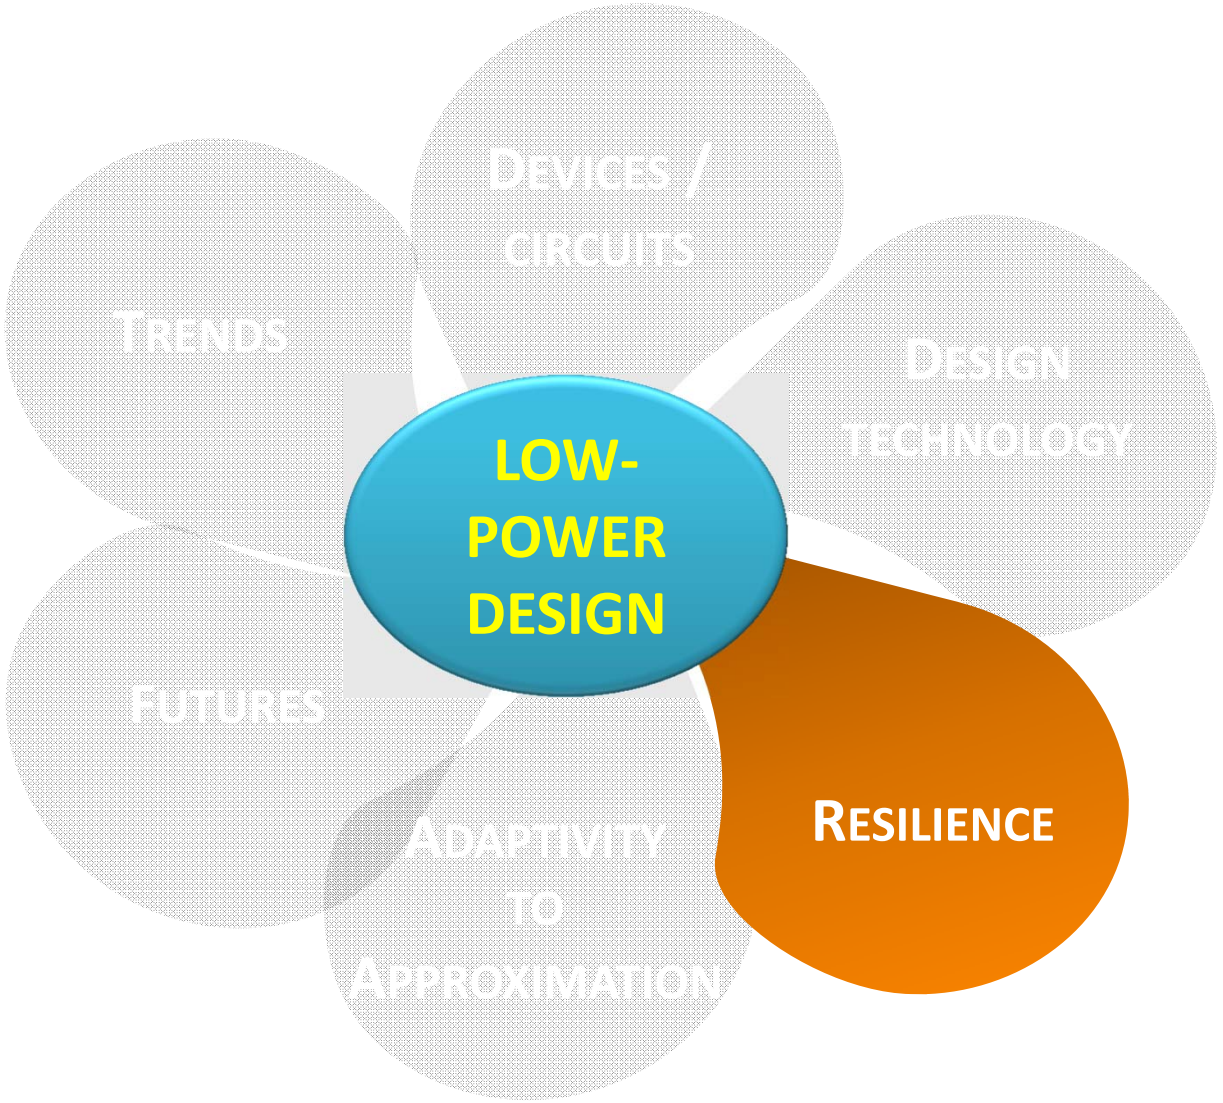
DEVICES / CIRCUITS
TRENDS
DESIGN TECHNOLOGY
LOW-POWER DESIGN
FUTURES
ADAPTIVITY TO APPROXIMATION
RESILIENCE

Resilience = "Long-Term Challenge" in ITRS

- **Resilience = system product's ability to mitigate variability, reliability phenomena**
- Error detection and repair mechanisms
- Alternative guardbanding mechanisms for different abstractions: stochastic, approximate, …
- "Cross-layer resilience" = recent buzz-phrase
- **Costs, benefits often hazy, difficult to quantify**

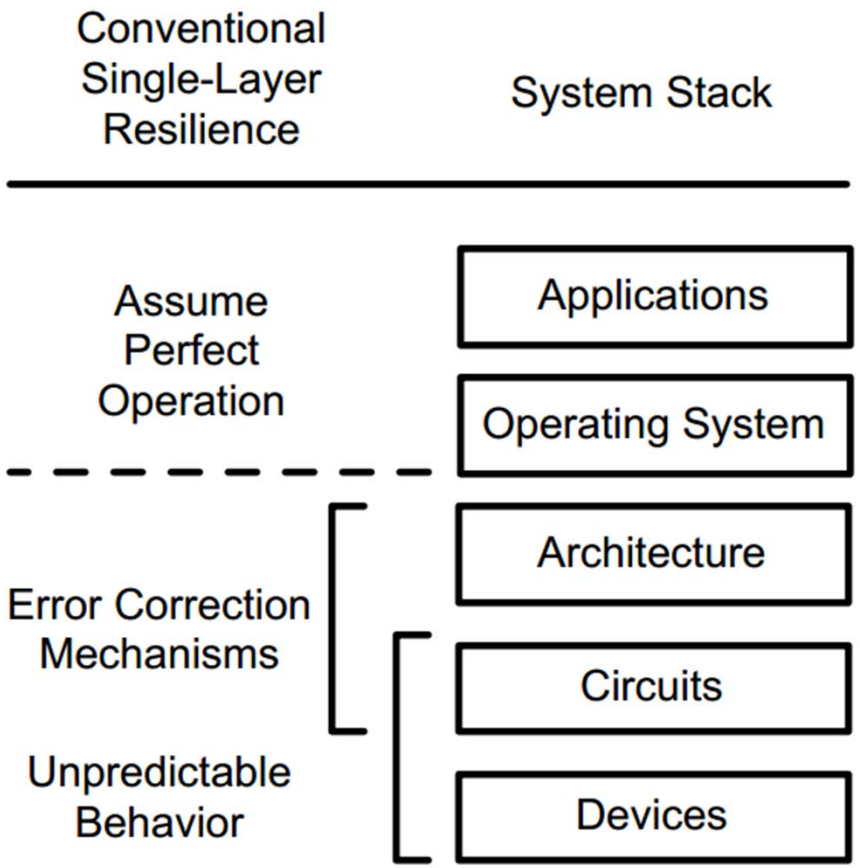

Example Step: Minimize Cost of Resilience

- Additional circuits ⇒ area and power penalties
- Recovery from errors ⇒ throughput degradation
- Large hold margin ⇒ short-path padding cost
- **Want benefits (e.g., energy) to maximally outweigh costs**

	Razor	Razor-Lite	TIMBER
Power penalty	30% [Das08]	~0% [Kim13]	100% [Choudhury09]
Area penalty	182% [Kim13]	33% [Kim13]	255% [Chen13]
#recovery cycles	5 [Wan09]	11 [Kim13]	0 [Choudhury09]

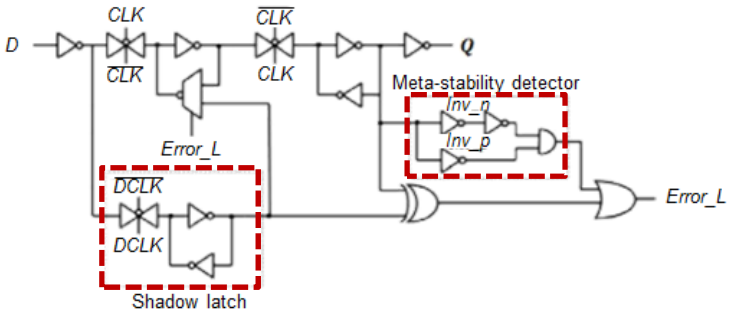

Razor

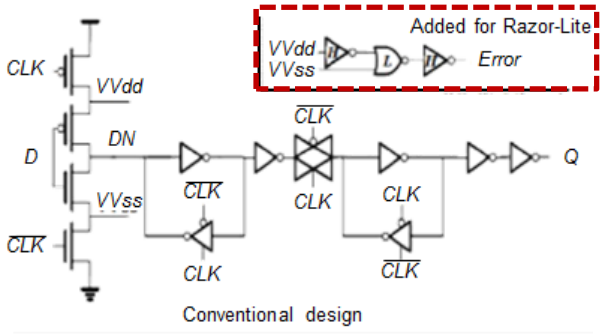

Razor-Lite

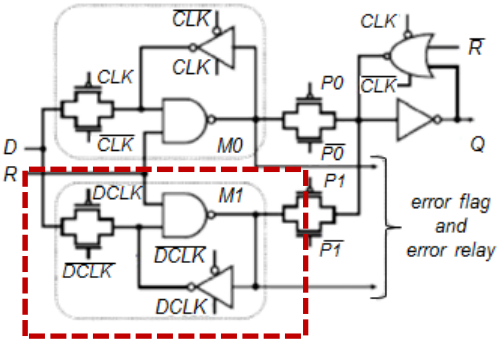

TIMBER

Tradeoff: Resilience Cost vs. Datapath Cost

#Razor FFs
(resilience cost)

Power/area of
fanin circuits

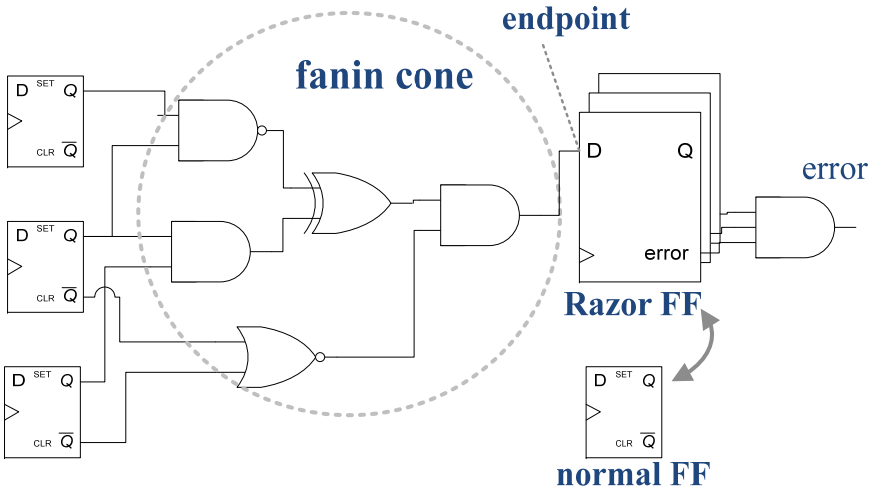

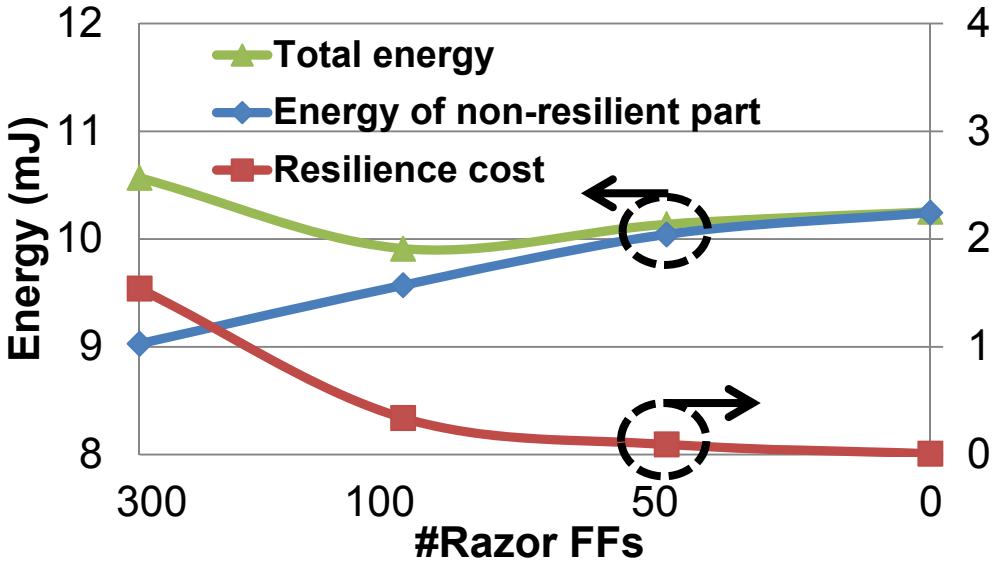

Can minimize total energy via this tradeoff

Selective-Endpoint Optimization (SEOpt)

- Optimize fanin cone of an endpoint w/ tighter constraints
  ⇒ Allows replacement of Razor FF w/ normal FF
- Pick endpoints based on heuristic sensitivity functions

Candidate Sensitivity Functions

$$SF1 = |slack(p)|$$

$$SF2 = |slack(p)| \times numcri(p)$$

$$SF3 = |slack(p)| \times \frac{num_{cri}(p)}{num_{total}(p)}$$

$$SF4 = |slack(p)| \times \sum_{c \epsilon fanin(p)} Pwr(c)$$

$$SF5 = \sum_{c \epsilon fanin(p)} |slack(c)| \times Pwr(c)$$

p = negative slack endpoint
c = cells within fanin cone
num_{cri} = number of negative-slack cells

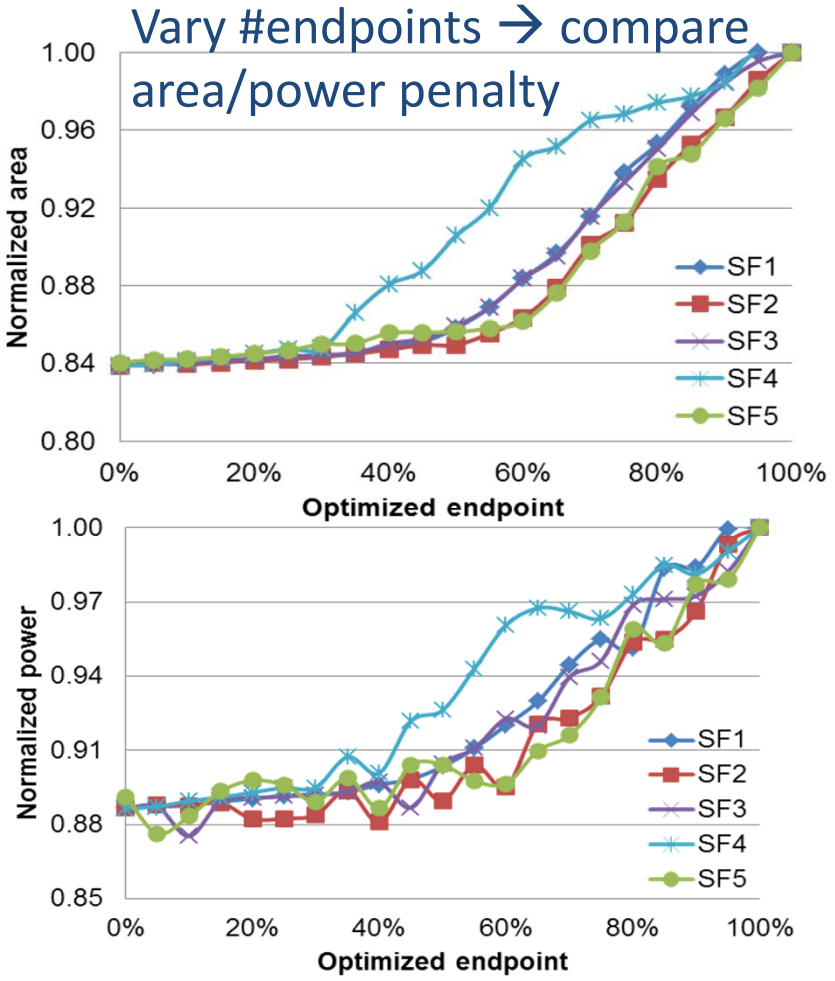

Clock Skew Optimization (SkewOpt)

- Increase slacks on timing-critical and/or frequently-exercised paths
  1. Generate sequential graph
  2. Find cycle of paths with minimum total weight
     → adjust clock latencies
     → contract the cycle into one vertex
  3. Iterate Step 2 until all endpoints are optimized

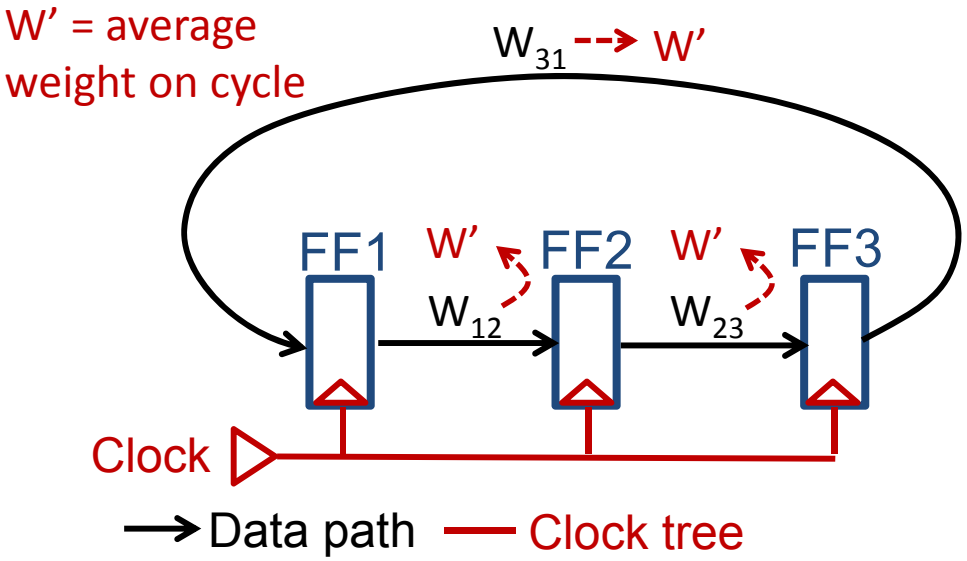

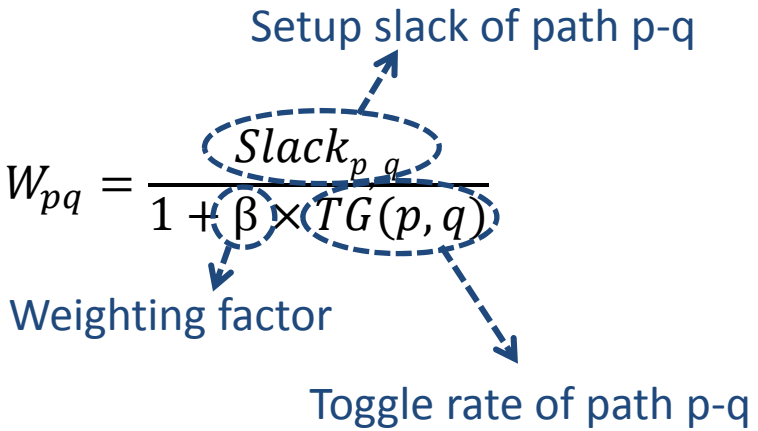

Overall Optimization Flow

[GLSVLSI14; ACM TODAES, to appear]

- Iteratively optimize with **SEOpt** and **SkewOpt**

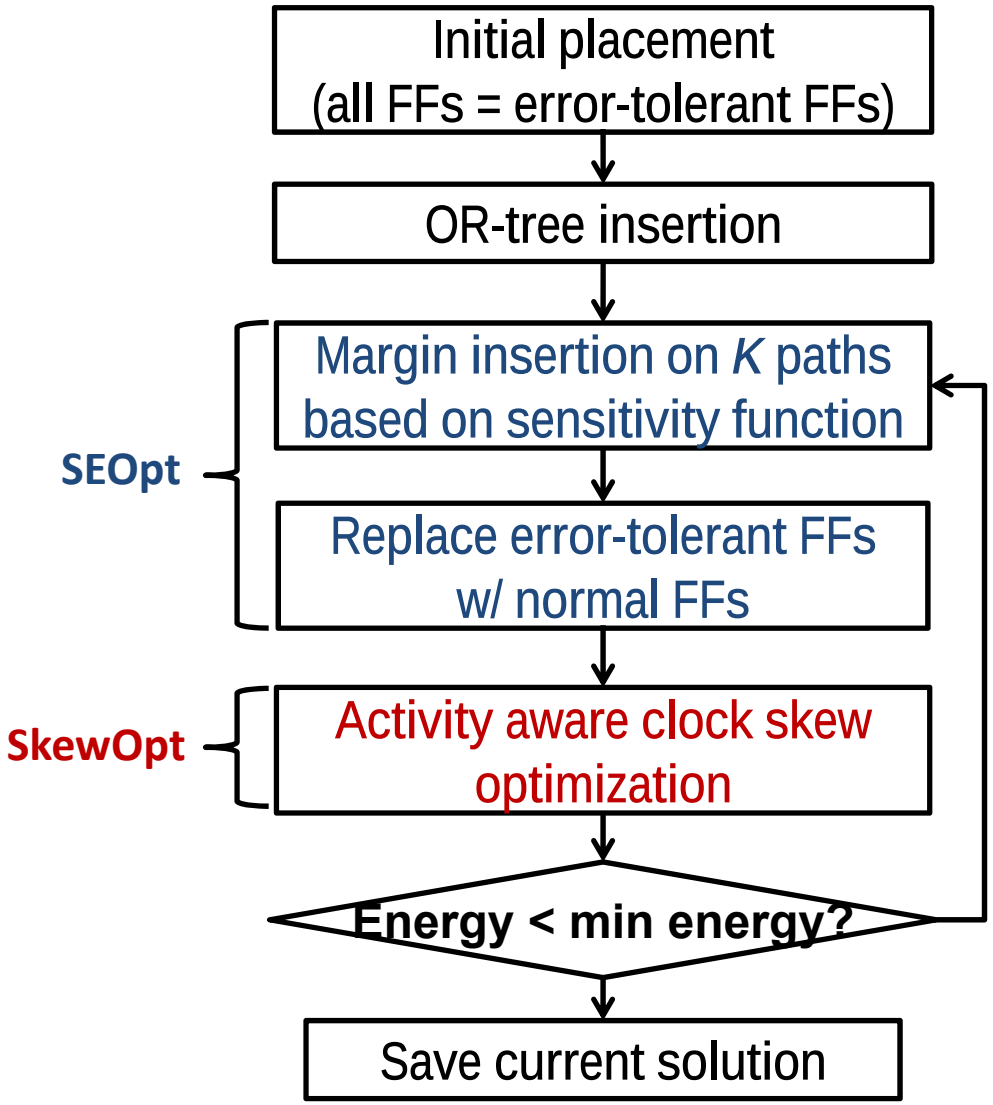

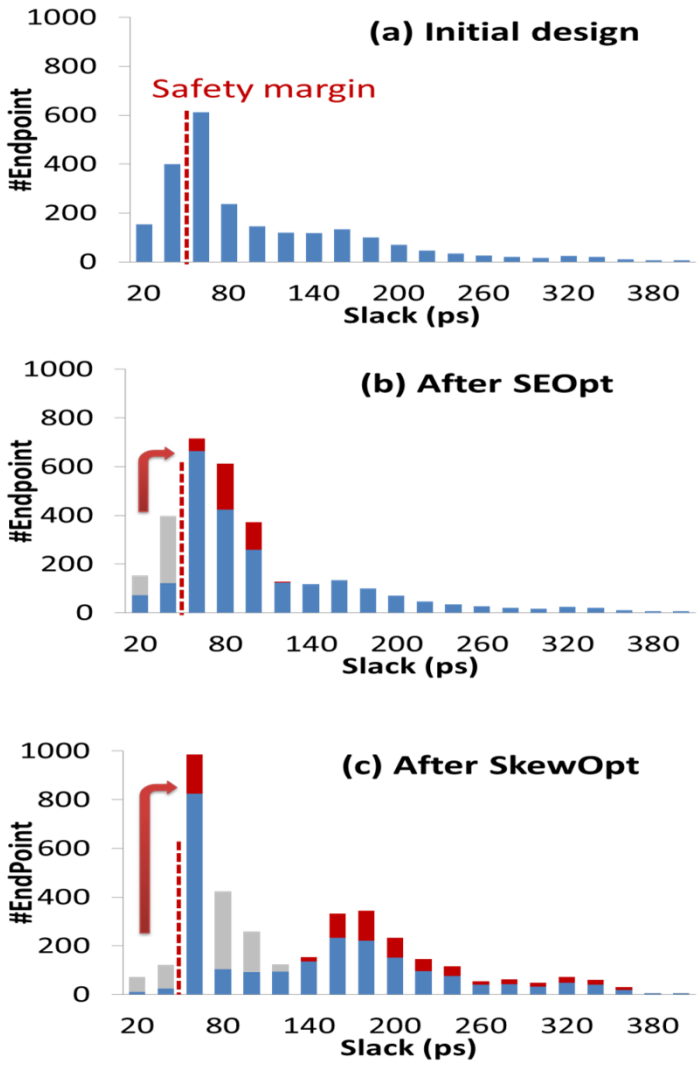

Benefit of Low-Cost Resilience

- Reference flows
  - Pure-margin (PM): conventional method w/ only margin insertion
  - Brute-force (BF): use error-tolerant FFs for timing-critical endpoints
- Proposed method (CO) achieves up to **21%** energy reduction compared to reference methods
- Resilience benefits increase with larger process variation

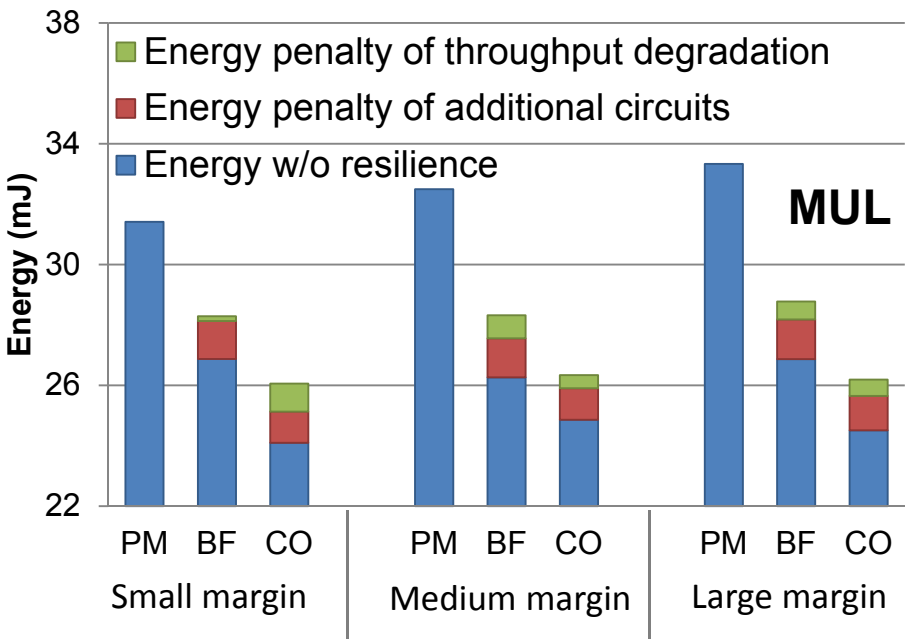

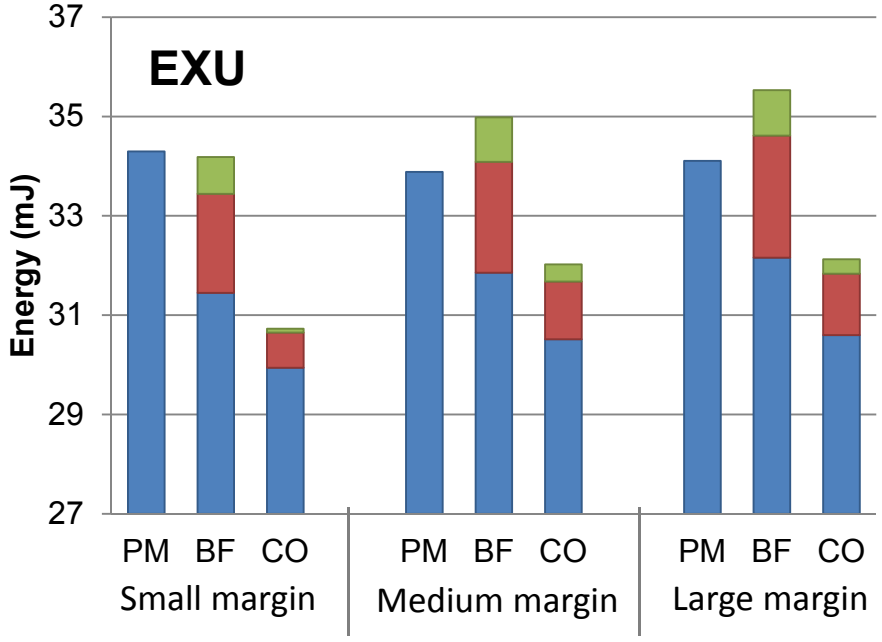

Small/medium/large margin ⇒ 1σ/2σ/3σ for SS corner

Technology: foundry 28nm

Increased Benefit of Resilience with AVS

- Adaptive voltage scaling allows a lower supply voltage for resilient designs, thus reduced power
- Proposed method trades off between timing-error penalty vs. reduced power at a lower supply voltage
- Proposed method achieves an average of 17% energy reduction compared to pure-margin designs
  ⇒ Resilience benefits increase in the context of AVS strategy

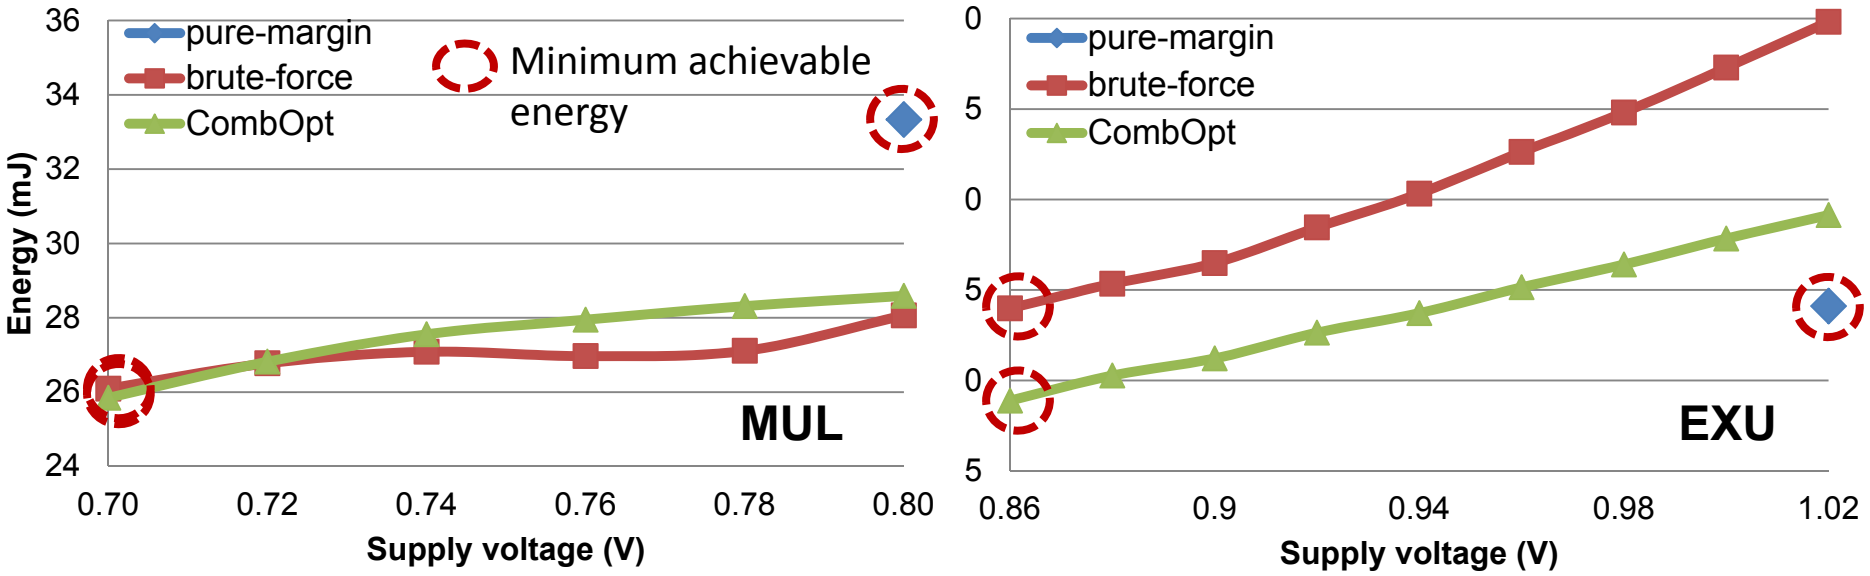

Technology: foundry 28nm

Agenda

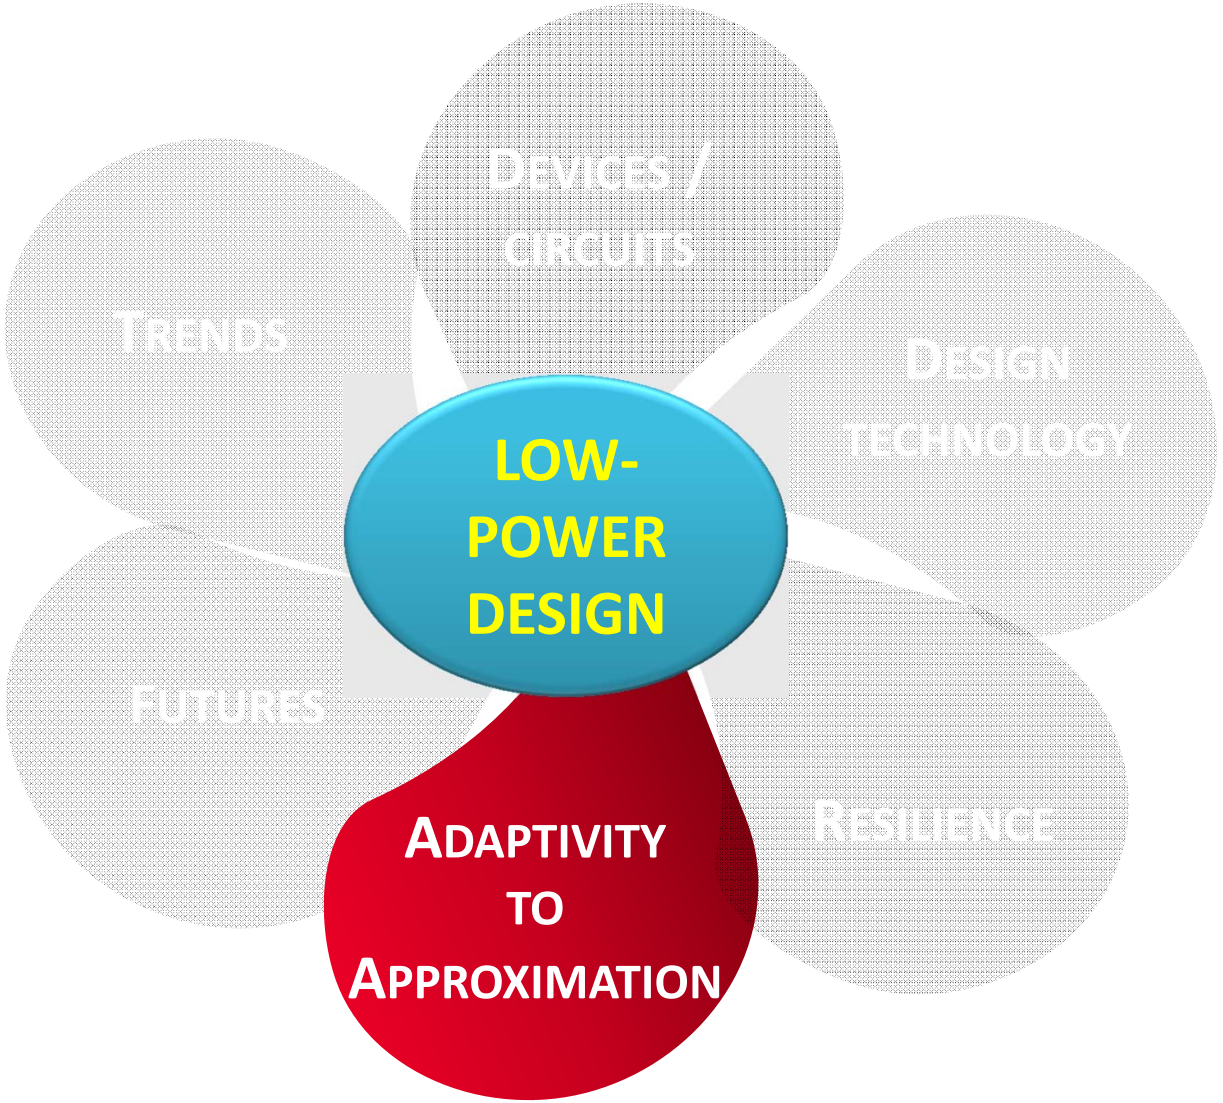
DEVICES / CIRCUITS
TRENDS
DESIGN TECHNOLOGY
LOW-POWER DESIGN
FUTURES
RESILIENCE
ADAPTIVITY TO APPROXIMATION

A Story of Adaptivity…

Step 1. Design team signs off chip at 1.4 GHz with worst-case (slow silicon) timing corner

Step 2. Chip comes back from fab (typical silicon) and runs at 1.8GHz

*[**Management is unhappy:** pessimism in signoff wasted area, power and design time]*

*[**Cross-functional tiger teams** are formed to work on (A) signoff corner pessimism (margin) reduction and (B) improved model-hardware correlation]*

Step 3. Scale down supply voltage so chip runs at 1.4 GHz with as little power as possible (= ***adaptive voltage scaling***)

Adaptive Voltage Scaling Approaches

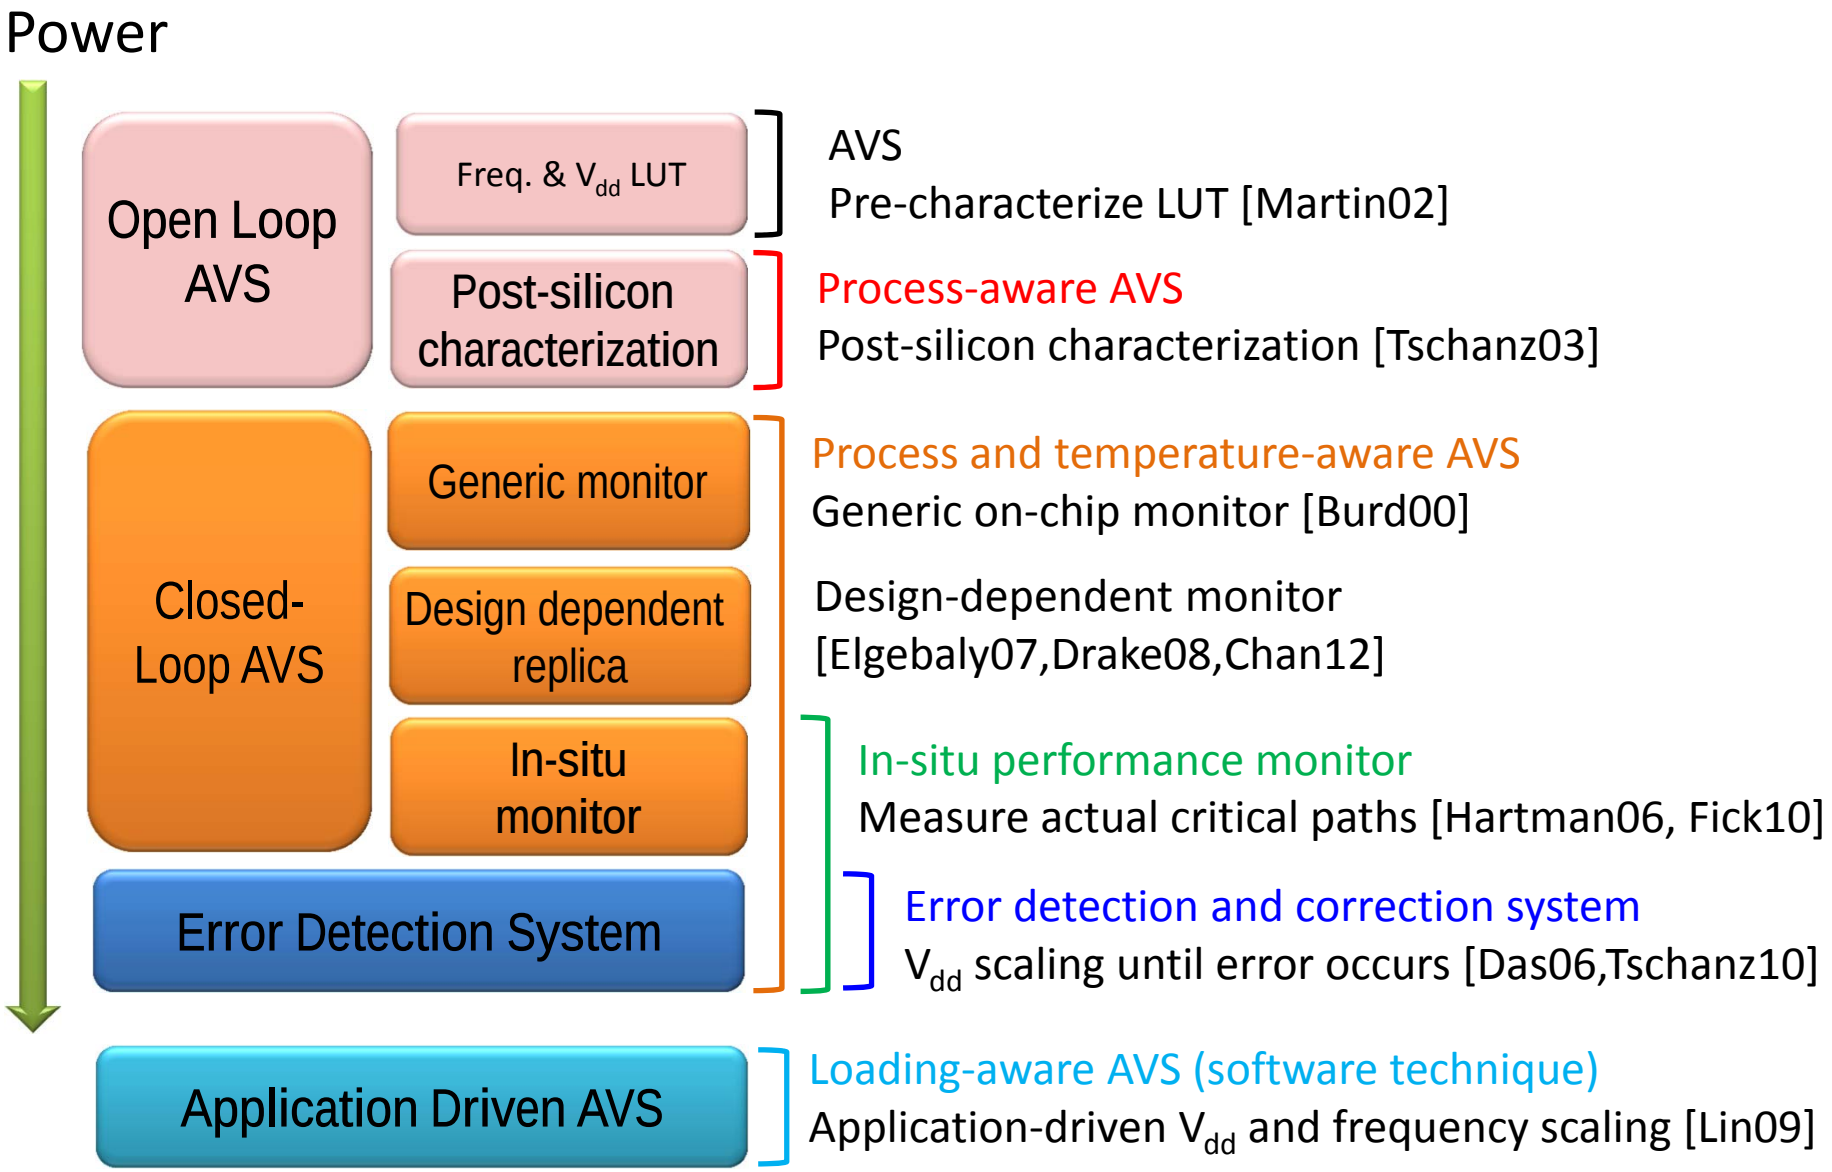

"Process-aware Voltage Scaling" (PVS) [ICCAD-2012]

- Monitor design considerations
  - Critical path can be difficult to identify (IP from 3rd party)
  - Multiple modes/voltages: F_{max} calibration requires long test time
- **UCSD : generic, tunable monitor**
  - **RO-based monitor with $V_{min\_ro} > V_{min}$ for any data path at any process condition (generic ⇔ overdesign)**
  - **Monitor is tunable based on F_{max} of sample chips to recover design margin (calibrate once)**
- Abstracts voltage scaling property instead of matching critical path
  - Keys: (1) PMOS-, NMOS-dominated paths determine V_{min}; (2) tune ROs with series resistance (pass gates)

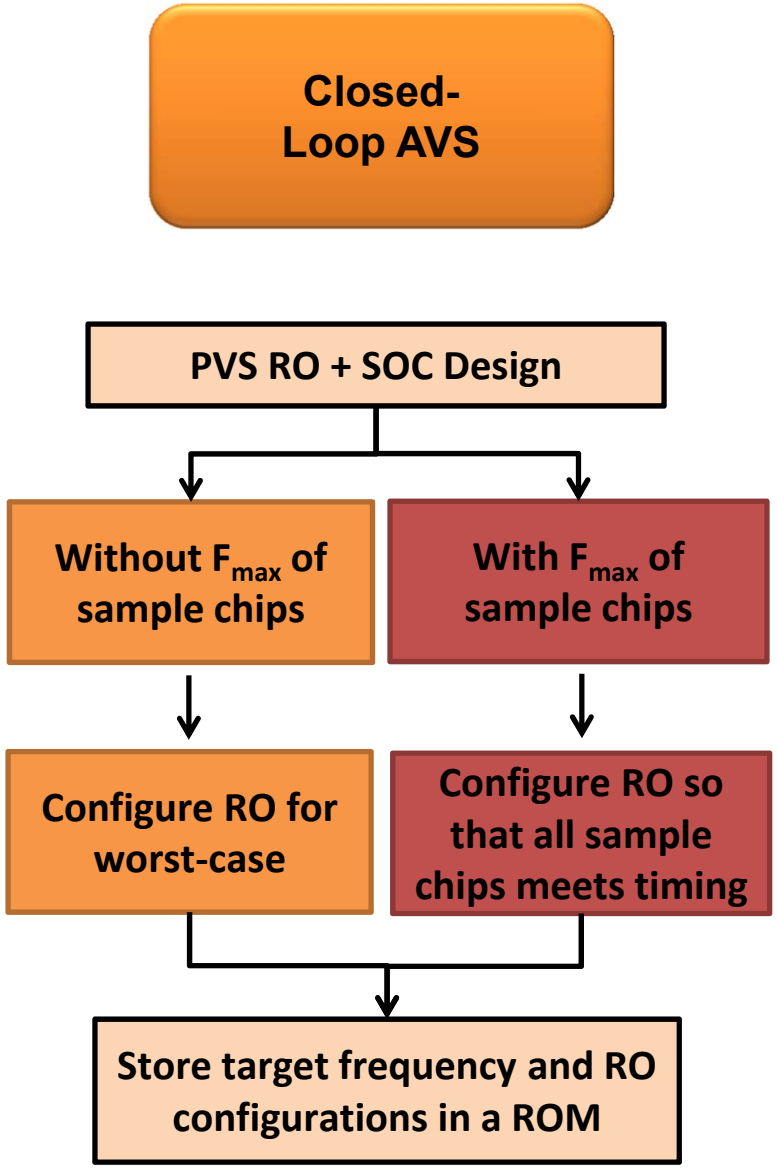

Voltage Scaling Basic Concepts

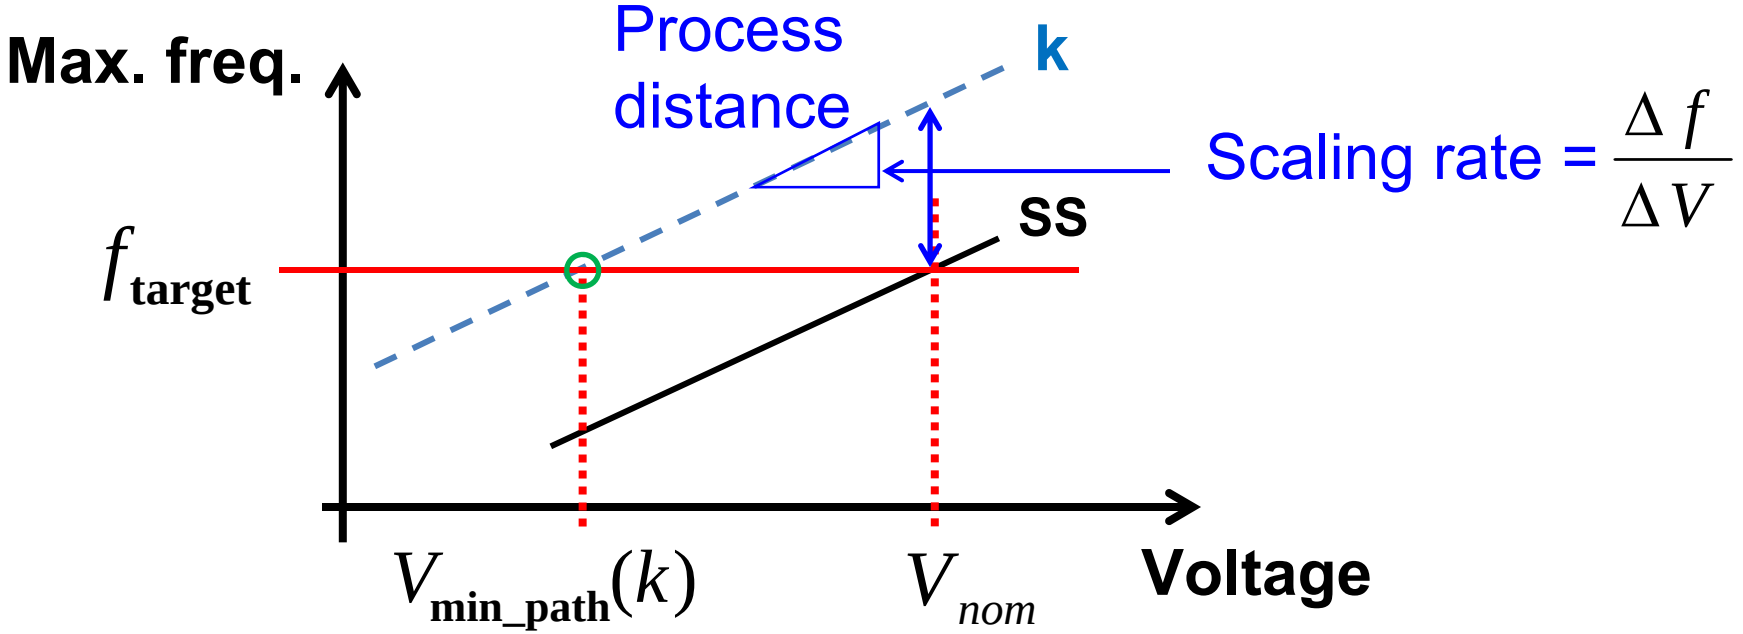

- Process distance: process-induced frequency shift relative to target frequency
- Scaling rate: frequency shift (Δf) per unit voltage difference (ΔV)
- V_{min}= Minimum V_{dd} to meet target frequency
  - Calculated from process distance and scaling rate

Process-Aware Voltage Scaling Concept

- Use V_{min} of ring-oscillator (RO) as a reference
- Design ROs with worst-case voltage scaling properties → an arbitrary circuit will meet target frequency at V_{min_ro}

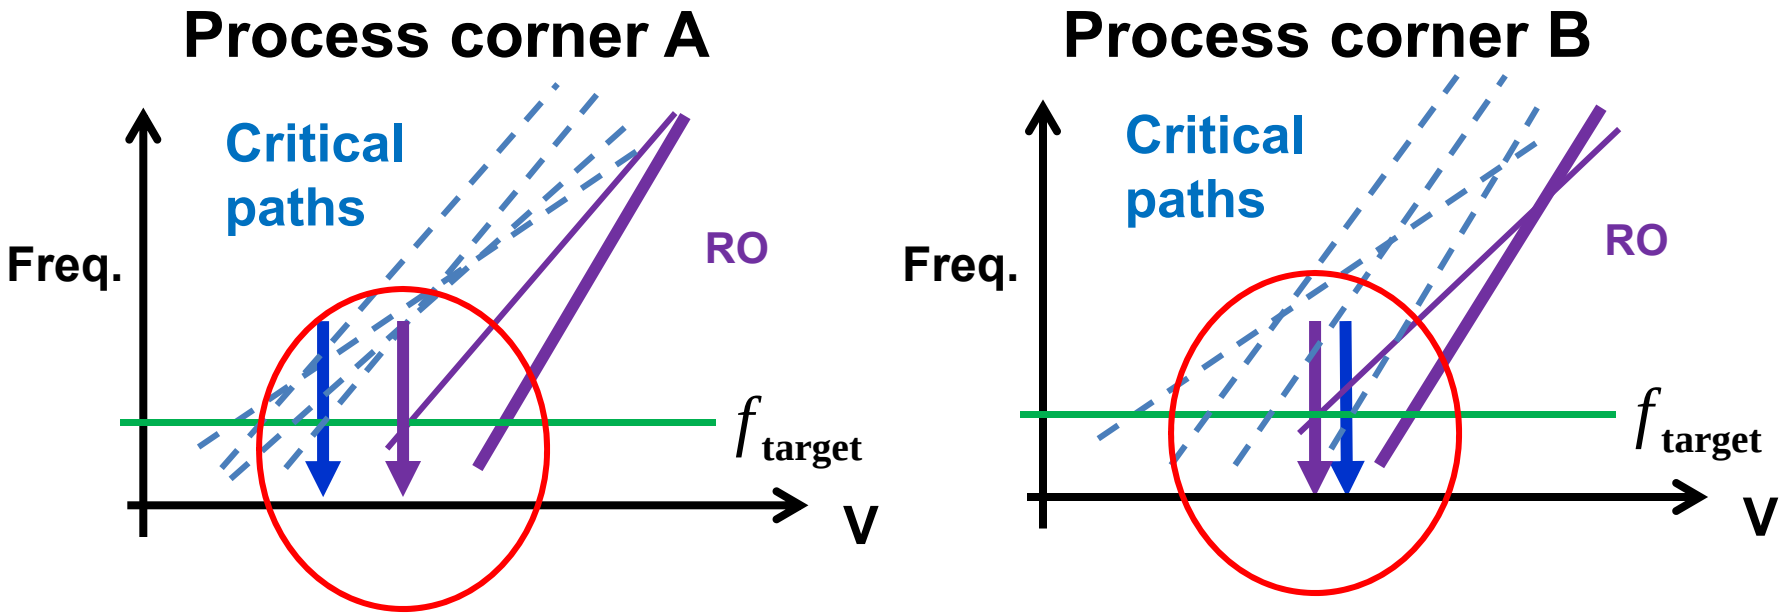

Max. V_{min} of ROs > Max. V_{min} of paths

Experimental Results on Tunability

Aggressive config.
→ $V_{min\_est} < V_{min\_chip}$
→ Some chips will fail

Optimized config.
- Increase % high resistance passgates
- $V_{min\_est} \approx V_{min\_chip}$

Default config.
- Low resistance passgates
- Guardband for worst-case
- $V_{min\_est} > V_{min\_chip}$
- 13mV margin

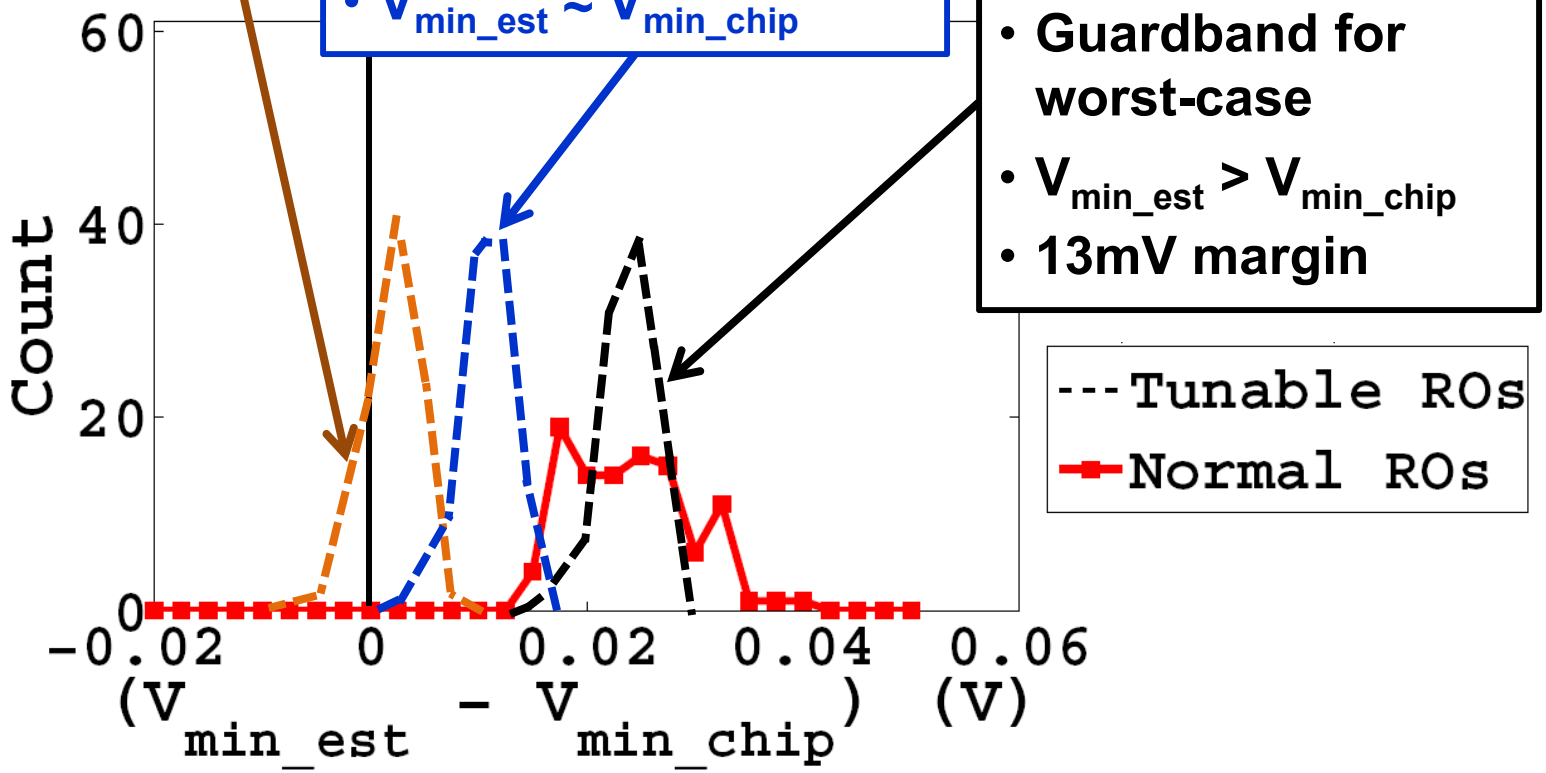

Experimental Results on Tunability

Aggressive config.
→ $V_{min\_est} < V_{min\_chip}$
→ Some chips will fail

Optimized config.
- Increase % high resistance passgates
- $V_{min\_est} \approx V_{min\_chip}$

Default config.
- Low resistance passgates
- Guardband for worst-case
- $V_{min\_est} > V_{min\_chip}$
- 13mV margin

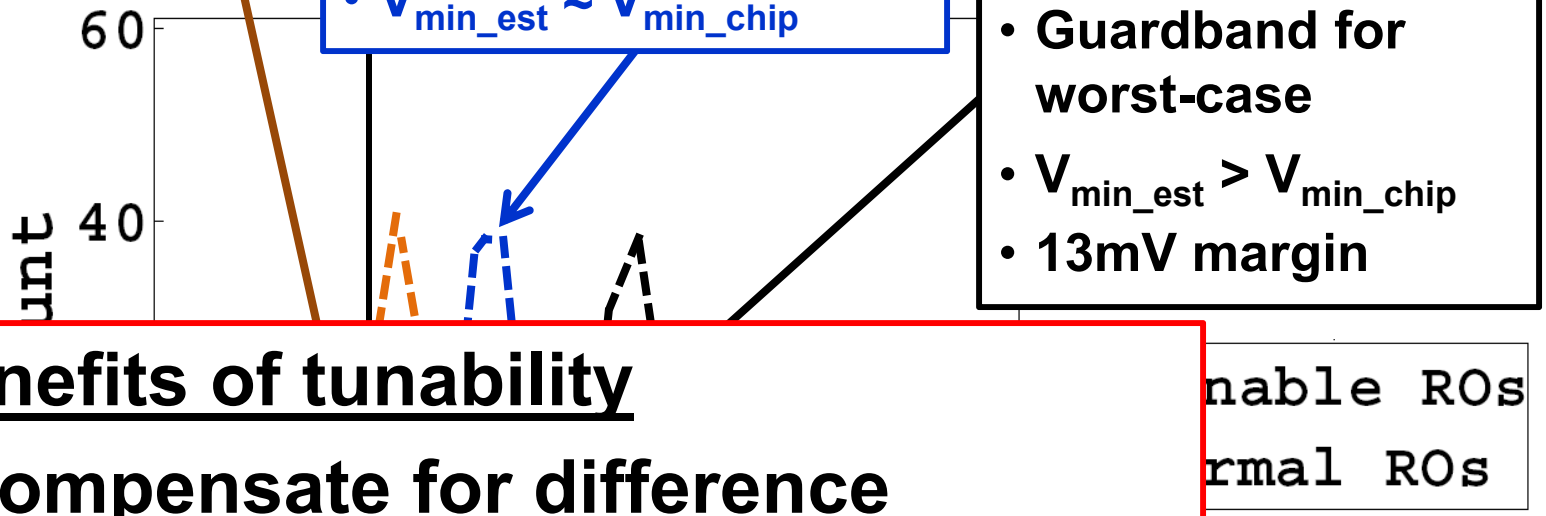

Benefits of tunability

- Compensate for difference between model vs. silicon
- Recover margin when variation is reduced due to improved process

Error-Resilience

- Error-freeness guarantees cost {margin, power, $$$}
- Unnecessary in some contexts
  - Machine learning, data mining, search
  - Signal processing: image, video, speech
  - Optimization
- Paradigms for error-resilience
  - Approximate computing
  - Stochastic computing
  - Probabilistic computing
  - …
- **In what contexts, with what knowledge ?**

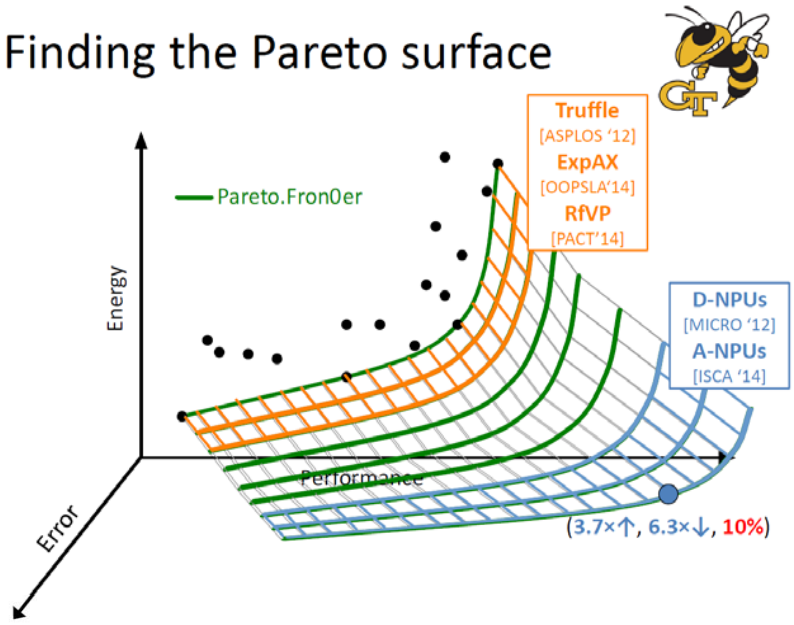

H. Esmaeilzadeh, "Approximation: Beyond the Tyranny of Digital Computing", IEEE RCS 2, May 2014

J. Han, M. Orshansky, "Approximate Computing: An Emerging Paradigm For Energy-Efficient Design", in Proc. ETS, 2013, pp. 1-6.

What If We Knew…(switching activity, workload)

Error-Tolerant Design

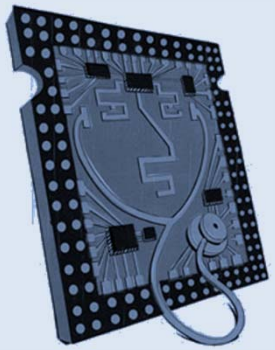

CPU, heal thyself …

Errors are detected and corrected with redundancy technique

Problem:

- Many paths have near-critical slack → *wall of (critical) slack*
- Scaling beyond the critical operating point causes massive errors that cannot be corrected

Reshape slack distribution for gracefully increasing error rate

Frequently-exercised paths
: upsize cells
Rarely-exercised paths
: downsize cells

→ **Scale voltage further**

Recovery-Driven Design for Error-Tolerance [TCAD12]

- Minimize power for a target error rate
- Slack redistribution based on functional information

Voltage Scaling

reduce voltage until the error rate exceeds a target

Path Optimization

optimize frequently exercised, negative slack paths

Power Reduction

reduce power without affecting error rate

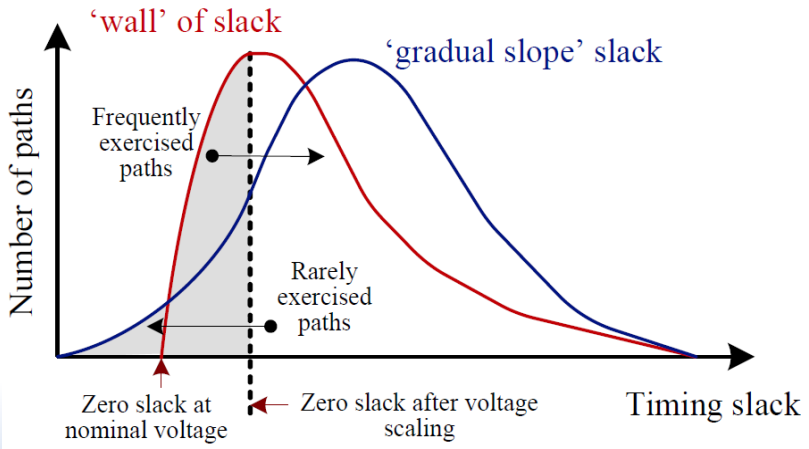

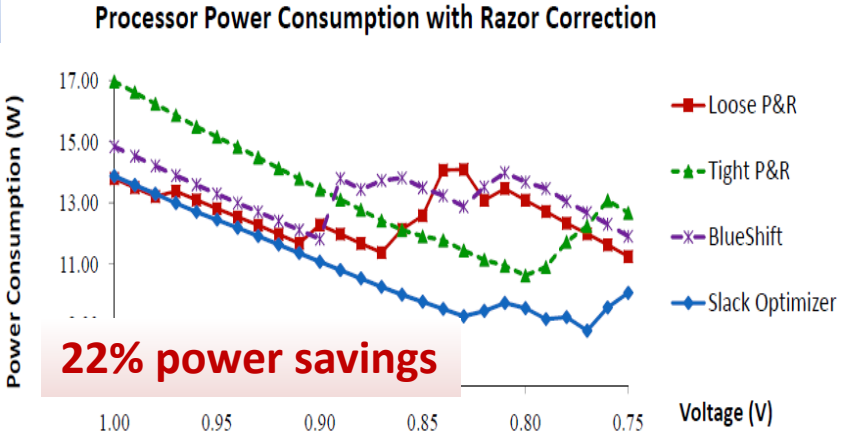

What If We Knew … (accuracy requirements)

Approximate Design

What is the square root of 10 ?

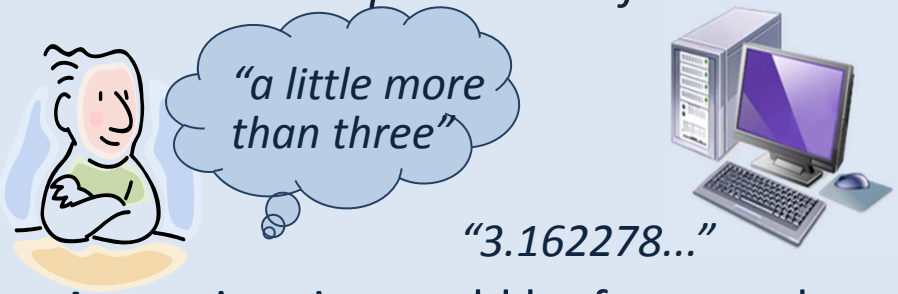

Approximation could be faster and more powerful

Problem:

- Accuracy requirement can change during runtime → benefits of approximation could be reduced

Adapt to changing requirements with runtime accuracy configuration

[DAC 2012]
"accuracy-configurable approximate adder"

lower power consumption ← → higher accuracy

[DAC12]

Accuracy-Configurable Approximate Adder

- Accuracy configuration with pipelined adder

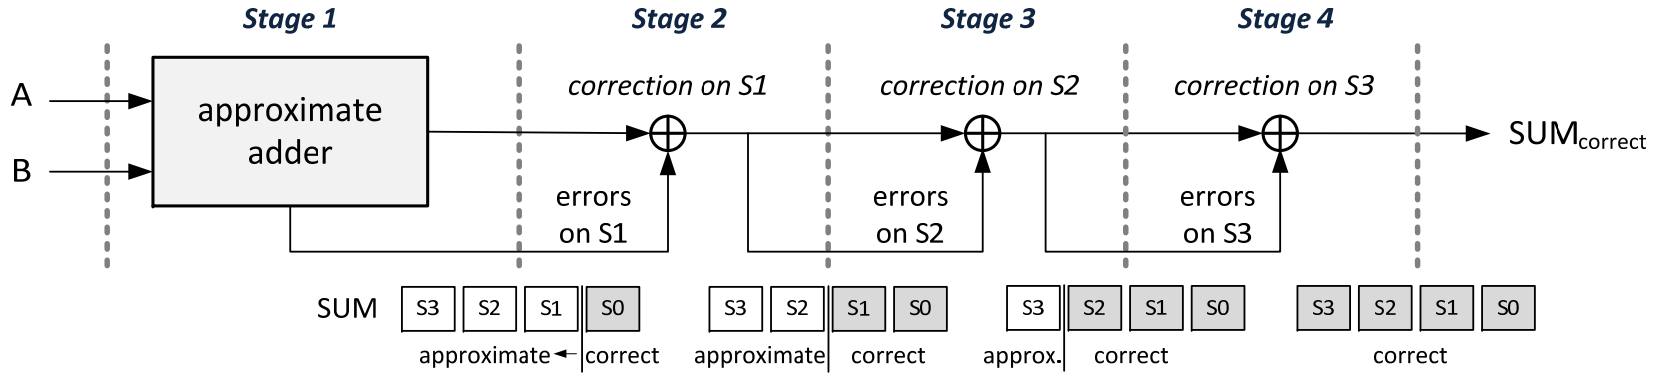

- Power reduction when accuracy requirement varying

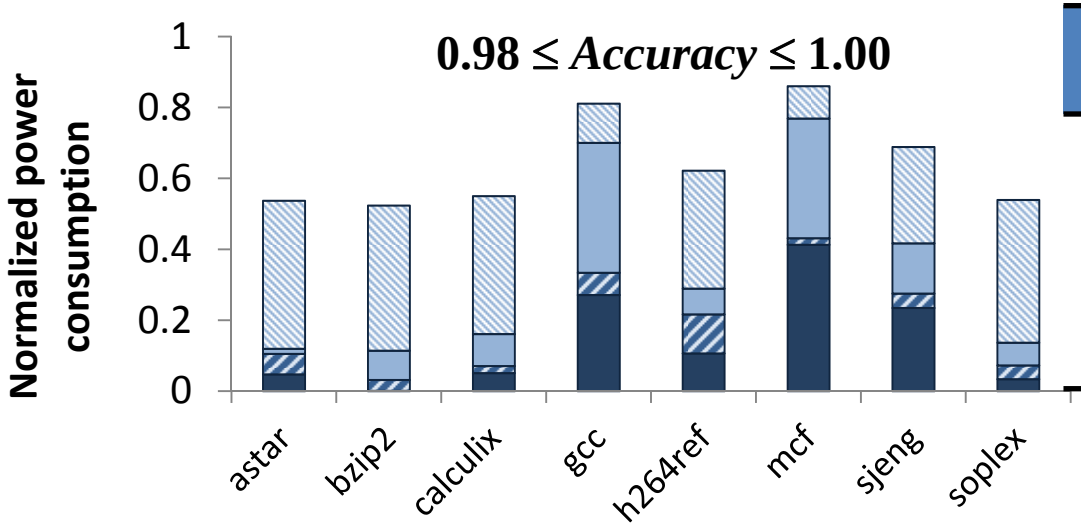

Config.	Accuracy	Power reduction
mode-4	1.000	-11.5%
mode-3	0.998	**12.4%**
mode-2	0.991	**31.0%**
mode-1	0.983	**51.6%**

$$Accuracy = Avg.\left(1 - \frac{|result - reference|}{reference}\right)$$

Average 30% power savings vs. no accuracy configuration

Stochastic Computing

- Conventional computation circuits rely on parallel bits aligned by clock → timing errors can be fatal (e.g., if occurring in MSB)
- Stochastic Circuit (SC) paradigm replaces parallel bits with serial bit-stream → resilient to voltage scaling

Example: Z= ¼+ ½ · X1·X2

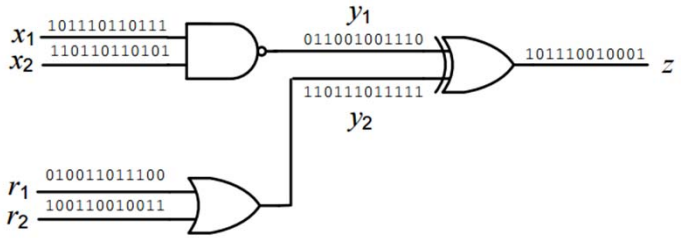

SC is resilient to voltage scaling

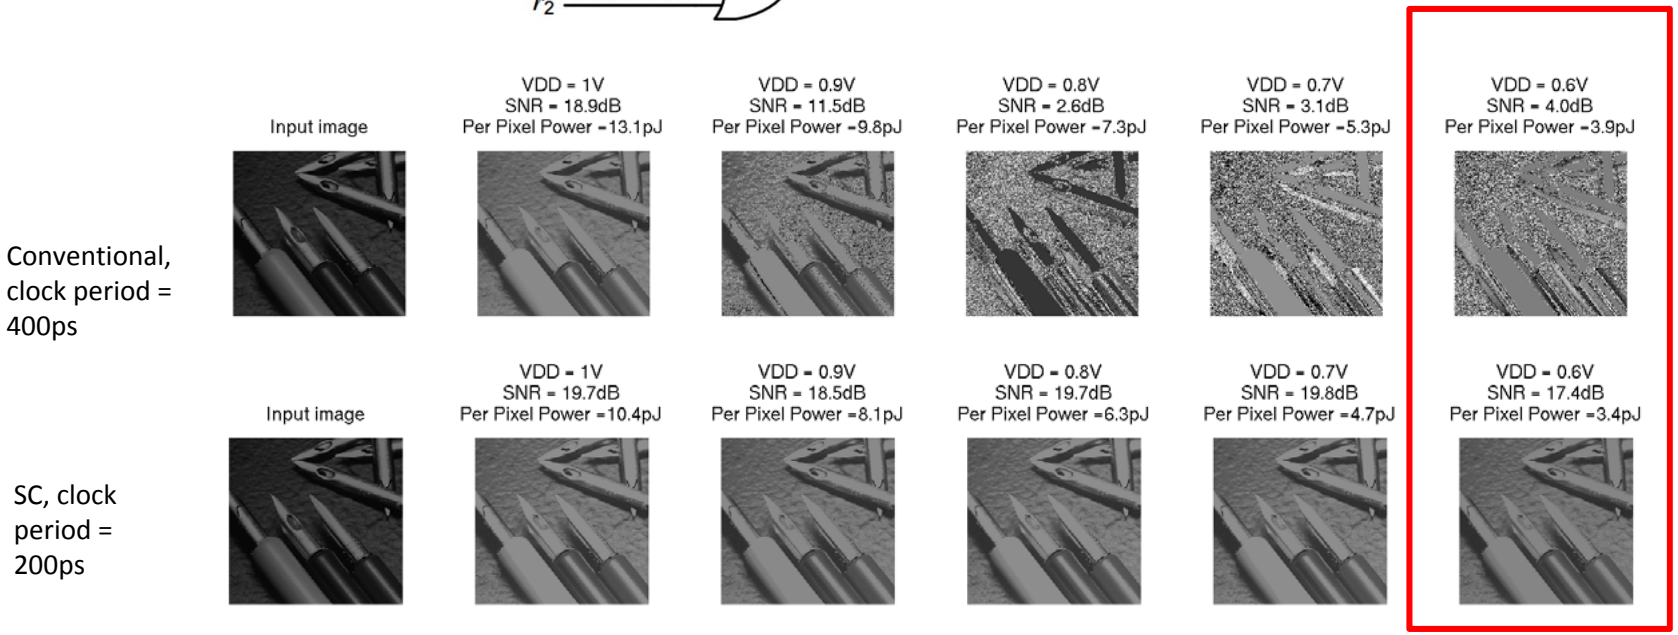

Agenda

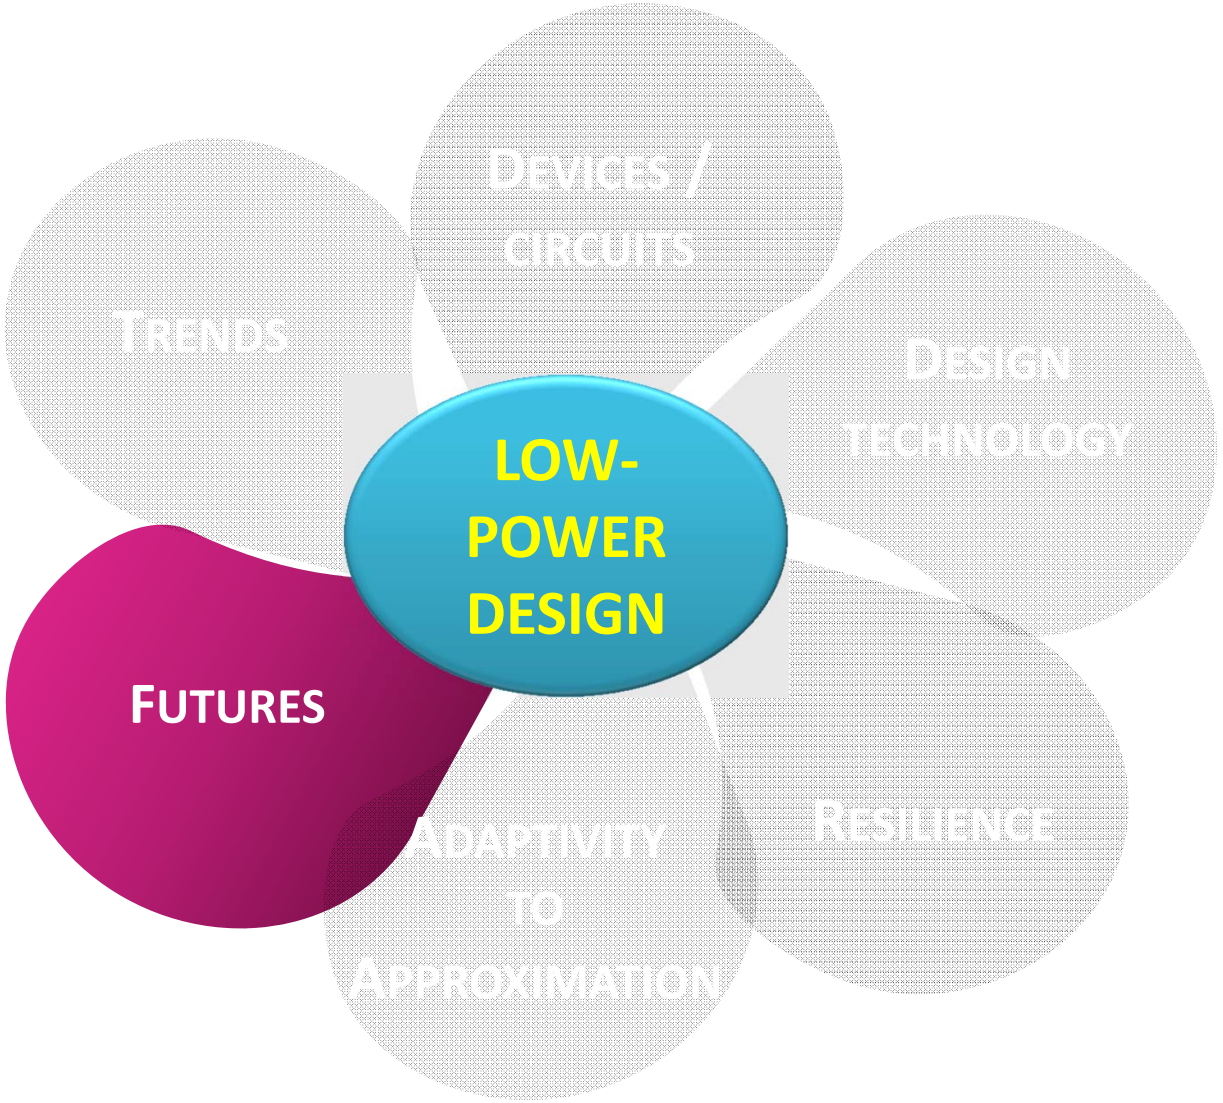
DEVICES / CIRCUITS
TRENDS
DESIGN TECHNOLOGY
LOW-POWER DESIGN
FUTURES
ADAPTIVITY TO APPROXIMATION
RESILIENCE

Better Optimizations

- Better MCMM gate sizing
  - Gate sizing with single timing view can induce timing violations in other timing view ⇒ multi-corner-multi-mode (MCMM) optimization is needed

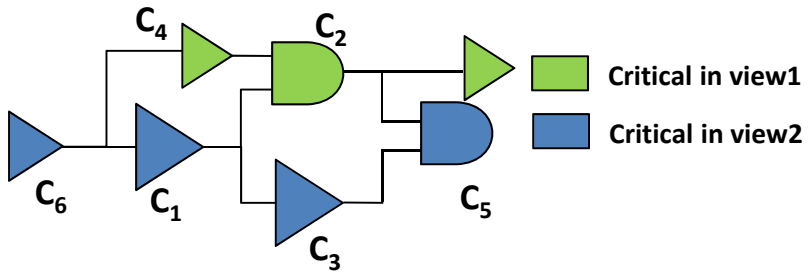

Upsizing C_2 for view 1 ⇒ Delay C_1 ↑ in view 2
⇒ Upsizing C_5 for view 2 ⇒ Delay C_2 ↑ in view 1
⇒ Upsizing C_4 for view 1 ⇒ Delay C_6 ↑ in view 2
⇒ ….

- Better FinFET fin discreteness-aware optimization
  - E.g., PlaceOpt to comprehend and avoid change in diffusion height induced by different number of fins

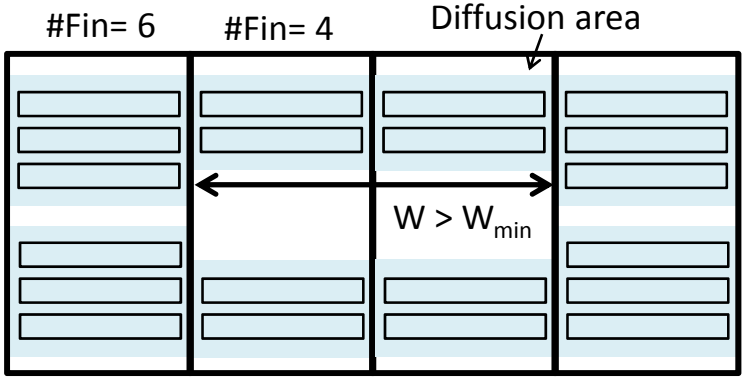

- Better design-technology co-optimization
  - E.g., BEOL stack optimizations, FEOL-BEOL gear ratio and library co-optimization for PPA, …

System Thinking: Design Synergy

- Low power implementation for the modern system on chip (SOC) requires a holistic and concurrent approach which includes collaboration between:
  - System level design
  - Architectural design
  - Software optimization and SW-HW co-design
  - Power aware RTL implementation/synthesis (front end design)
  - Physical design (chip/block level) (back end design)
  - IP design:
    Circuit design, Physical implementation of the IP
  - Process selection ((Bulk CMOS, SOI, FinFET) and device definition (nmos, pmos etc). Process optimization and DFM.
  - Adaptive design (on die sensors, process aware voltage scaling, process, power and temperature monitors). Power models (upf/cpf)
  - Power and thermal verification and modeling
  - Silicon characterization, power models validation, silicon to model correlation

S. Dobre, UCSD lecture, 2015.

Example: Power and Performance Meters

- Power and performance meters:
  - Hardware–software solution
  - Measures performance and power in real time for different sub-systems integrated in the system on chip
  - Provides feedback to the system for:
    - Power management
    - Thermal management
    - Workload optimization

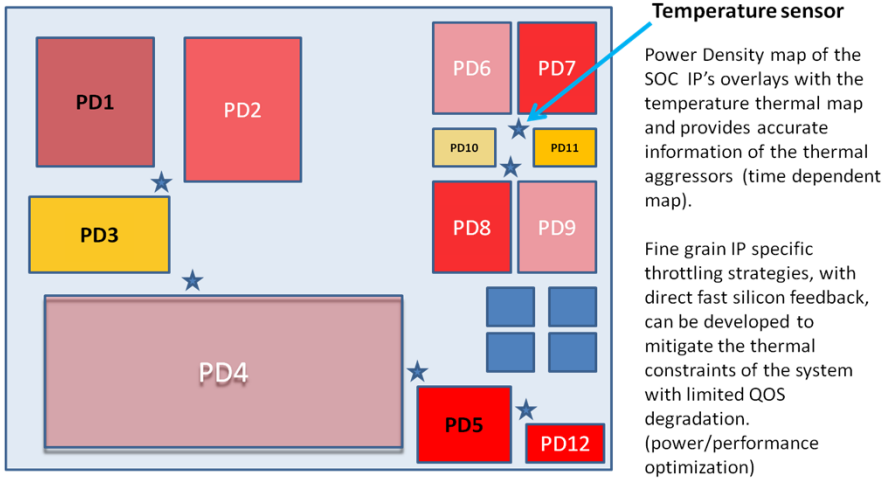

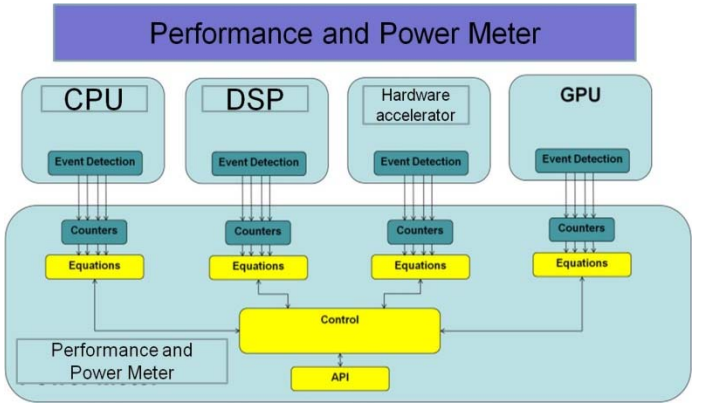

S. Dobre, UCSD lecture, 2015.

What If We Knew…(scenarios, duty cycles)

- DVFS allows adaptation to workloads, operating conditions
- DVFS processor operates at multiple power/performance points with different lifetimes
- **Lifetime energy can be different in each scenario (R * X)**

Different duty cycle (R)

Different frequency scaling (X)

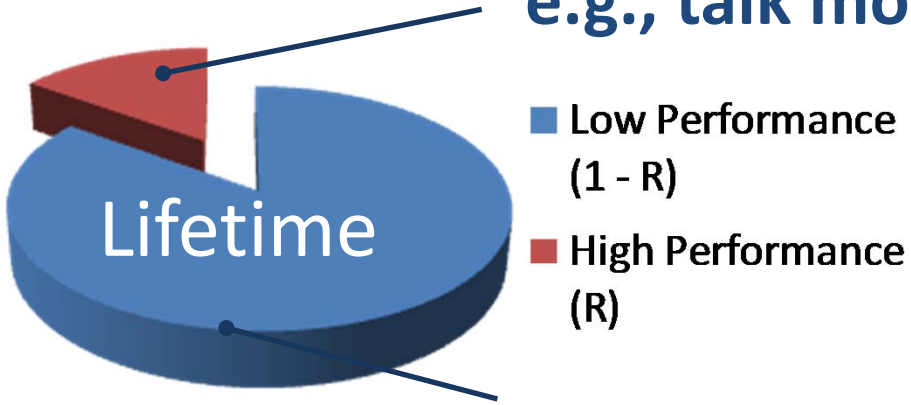

$$X = \frac{\text{clock frequency of high-perf mode}}{\text{clock frequency of low-perf mode}}$$

Context-Aware Multi-Mode Low-Power Design

Minimize lifetime energy based on modes, duty cycles

Multi-mode design

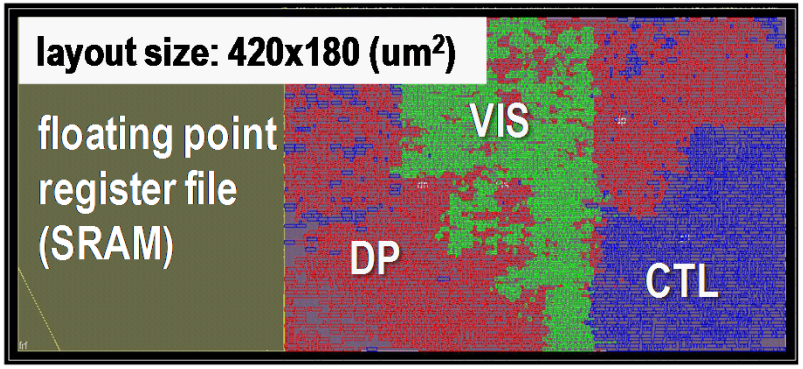

Selective-replication design

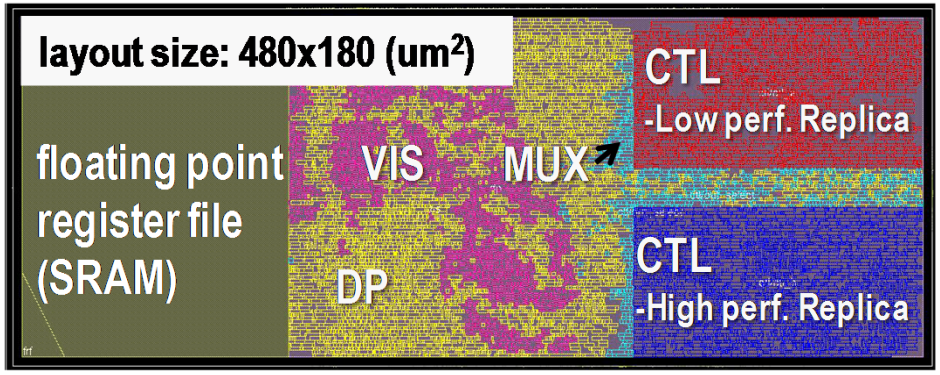

- CTL module has 12% energy savings through replication
- Processor-level: selective replication gives 12% total energy savings with 10% area overhead

Lower Power With 3DIC

- Power = key value proposition for 3DICs (shorter / wider connections)
- Recent work (DAC-2015): 3DIC power reduction at 28nm foundry FDSOI libraries and estimation of 3D power benefits only from 2D implementations
- Power benefit with 3D varies with testcases and implementation styles
  - Percentage delta benefit ranges from -5.1% (i.e., power increases in 3DIC) to 16.0%

3D structure: F2B and F2F

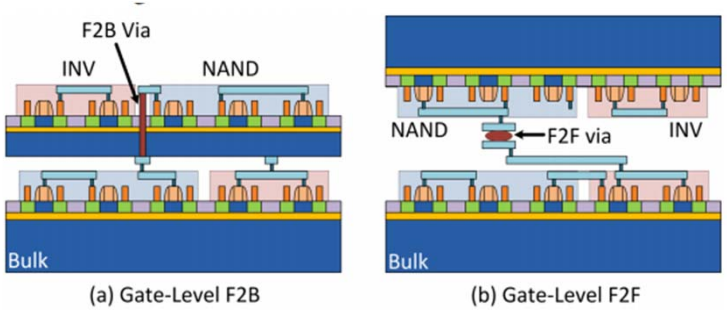

3DIC Implementation and Modeling flow

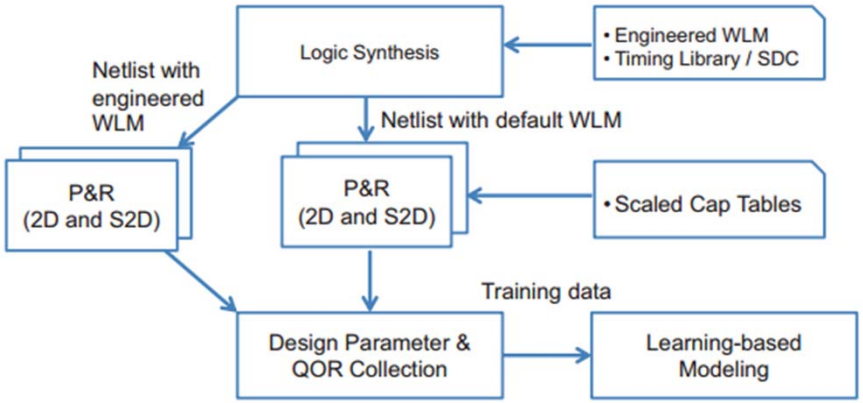

3DIC % delta power benefit relative to 2DIC Implementation

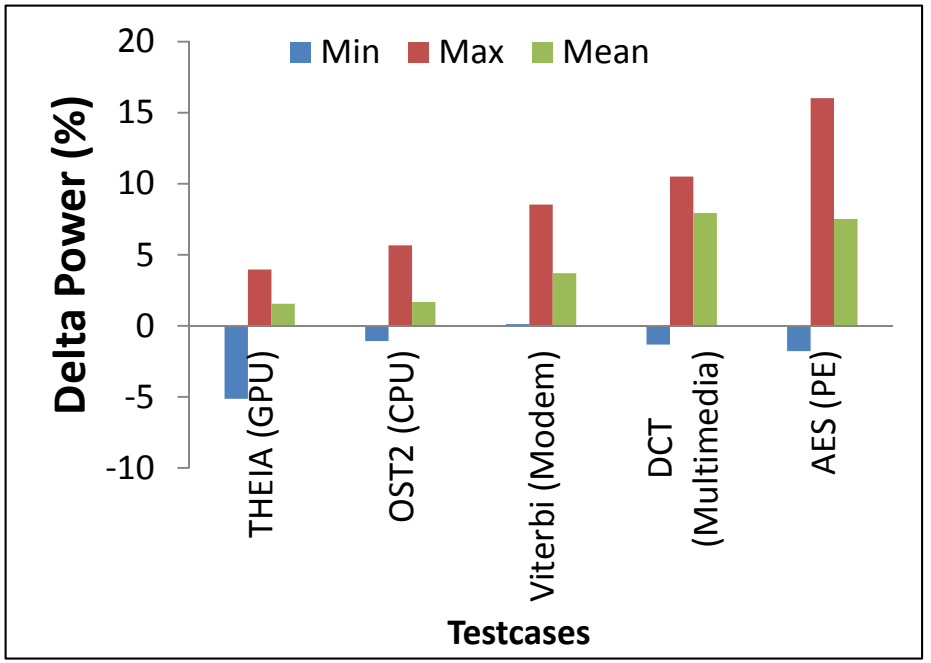

Clock (well-known, but still on table)

- Clocking = ~30-40% of total power
- SOC complexity: 1000+ clock domains
- Clock architecture: Many frequencies, local dividers to reduce frequency
- SP&R: planning, placement, buffering of “top-level”: CGCs, MUXes, dividers
- CTS: MCMM skew reduction, skew variation reduction (high voltage is wire-dominated but low voltage is gate-dominated, large hold buffering costs, …)
- Routing: NDRs, routing with distributed drivers, long common paths (reduces wirelength, driver sizes and number of clock buffers)

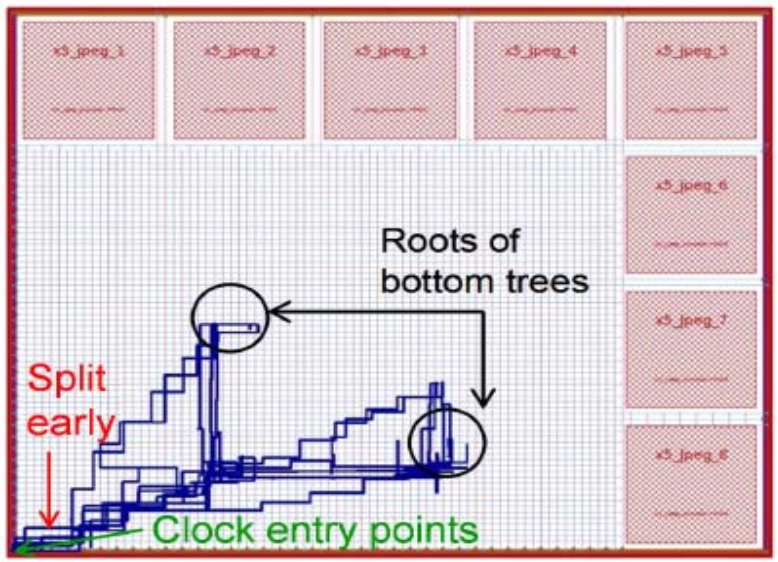

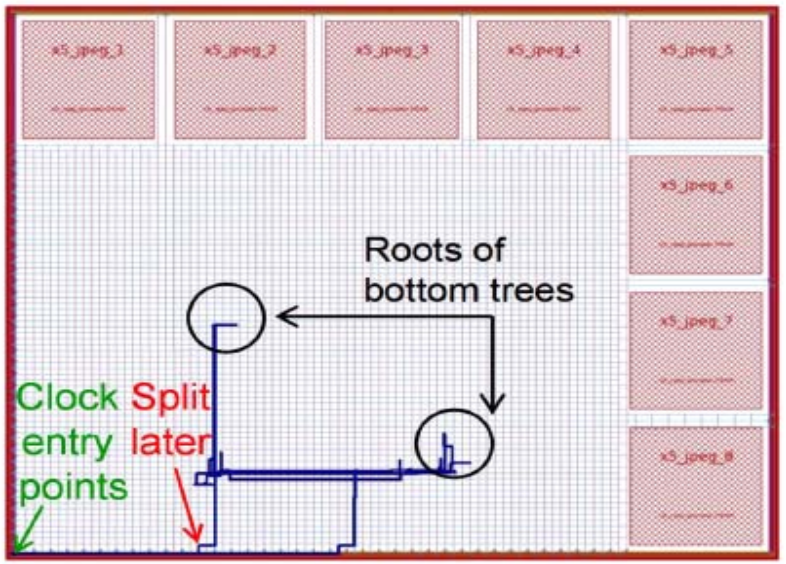

T.-B. Chan, K. Han, A. B. Kahng, J.-G. Lee and S. Nath, "OCV-Aware Top-Level Clock Tree Optimization", Proc. Great Lakes Symposium on VLSI, 2014, pp. 33-38.

Clock (Emerging)

- Reduced-swing, half-swing clocks
  - Reduces clock power when V_{tn} (resp., V_{tp}) is less (resp. greater) than ½ V_{dd}
  - On-off characteristics of NMOS (resp. PMOS) remain unchanged

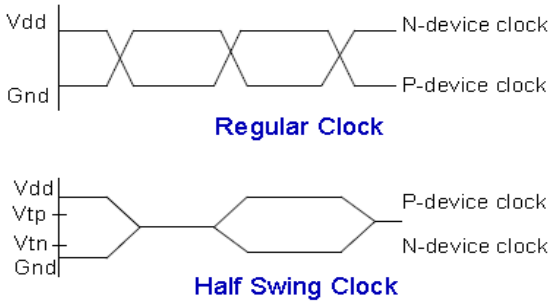

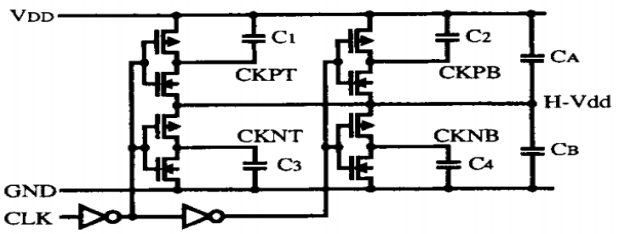

Half-Swing Clock [1]

- Globally Asynchronous, Locally Synchronous (GALS); **more asynchrony**

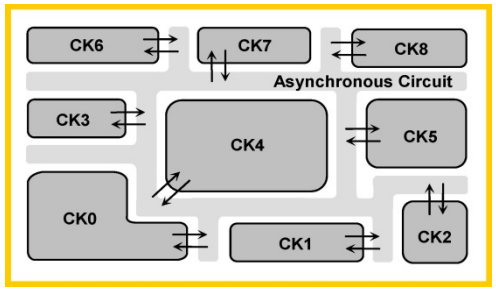

GALS Architecture [2]

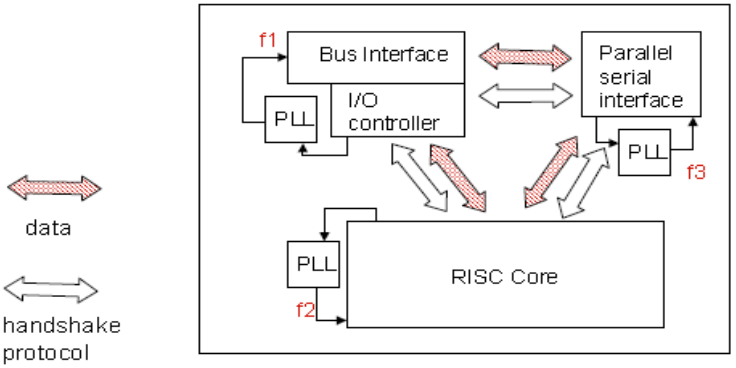

GALS with RISC Core [3]

Sources:

[1] H. K, S.Tanaka, K. Sasaki, "Half-Swing Clocking Scheme for 75% Power Saving in Clocking Circuitry", JSSC 30 (4) (1995), pp. 432-435
[2] I. Miro Panades, A. Greiner, A. Sheibanyrad, "A Low Cost Network-on-Chip with Guaranteed Service Well Suited to the GALS Approach", NanoNet, 2006
[3] http://www.cse.psu.edu/~mji/asic/tutorial-clock.pdf

A Few Nutshells…

- **AVS** mandatory to address variation at low voltage
- **NTC** in use as "ultra-low voltage mode"; standard design enablement works out of box in FinFET nodes (e.g., 0.46V – 1.25V)
- **FinFET transition** has large benefits → takes some pressure off of low-power design in near term (?)
- **Next-generation optimizers** needed for DVFS/MCMM, wide voltage corners, multi-patterning, FinFET discreteness, …
- **Next-generation analysis tools/flows** (thermal, dynamic IR, reliability wearout, stress = "next loops to close") to further squeeze design margin
- **Active power** regains focus → clock power reduction, active leakage cost, design for min V at given throughput
- **System-level low-power design** continues as both (complexity) challenge and (cross-layer) opportunity
- **3DIC** will turn the corner for power-performance envelope
- **Approximate computing** will take longer to turn the corner …


THANK
YOU!

Backup

Power = Gorilla + Elephant …

Older techniques

Domain	Technique	PDyn	PStat
Arch	Multilevel caches		+
	Multithreading		+
	Hardware virtualization		+
	Superscalar		++
	SMP	+	
Circuit	Low power physical libraries	+	+
	Back biasing		+
	Adaptive body biasing	+	++
	Power gating	(-)	++
	DVFS	+	

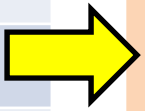

Newer techniques

Domain	Technique	PDyn	PStat
SW	Virtual prototype	+	+
	Many-core dev tools	+	
	Power-aware	+	
HW-SW	Co-partitioning at the behavioral level	+	
Arch	Heterogeneous parallel processing	+	
Circuit	Frequency islands	+	
	Near-threshold	+	(-)
	Asynchronous	+	

Source: 2011 ITRS (www.itrs.net/Links/2011ITRS/Home2011.htm)

Note: Guardband = Overdesign = Power

Design margin = stack of layers of conservatism

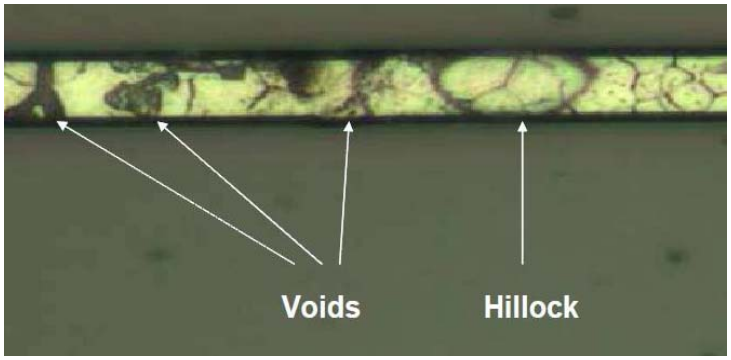

Reliability

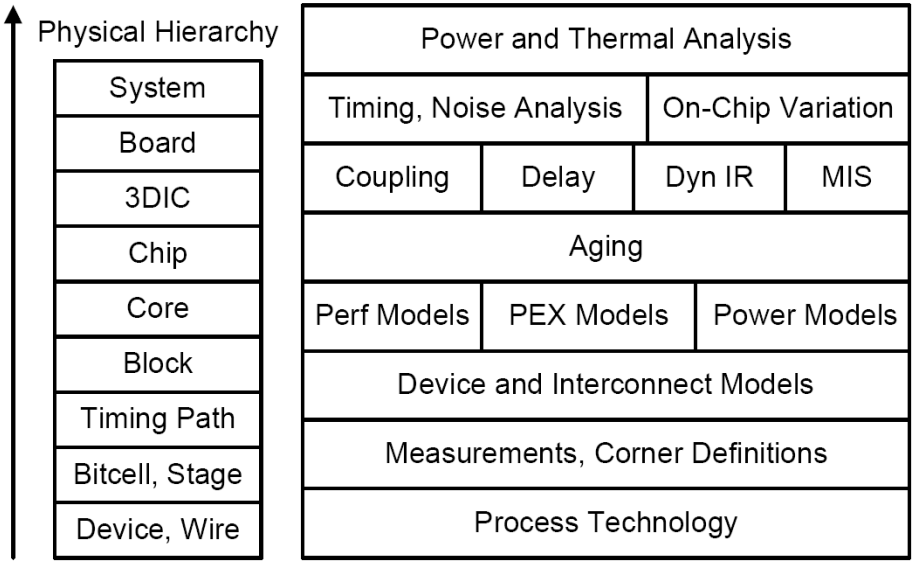

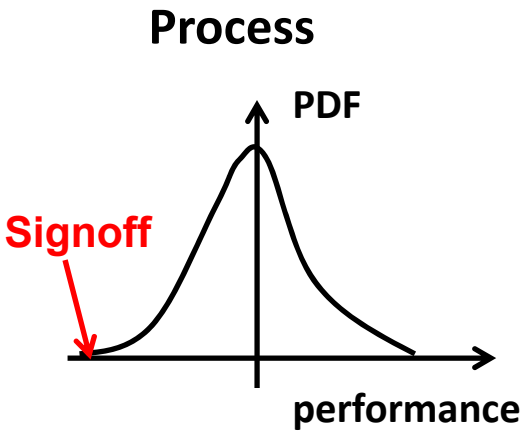

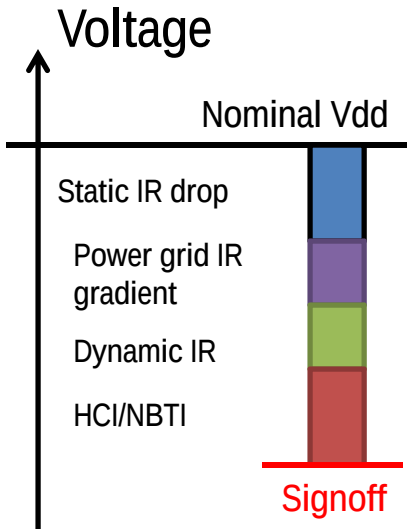

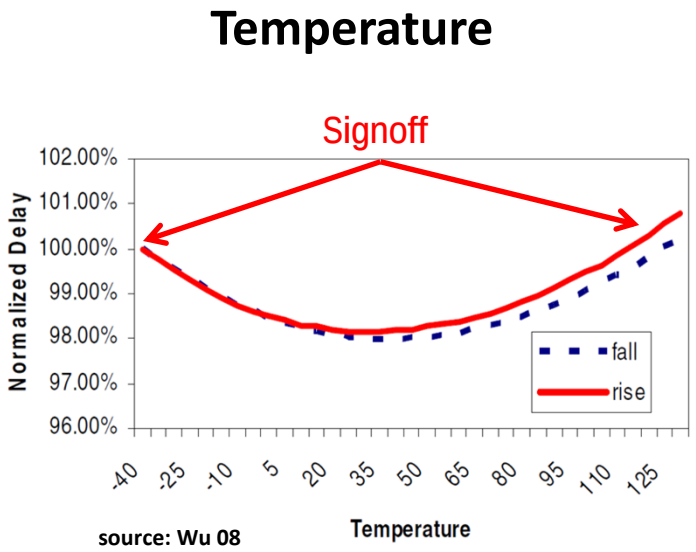

source: Wu 08

2008 Study: Cost of Guardband

ISQED08

Question: What is the concrete benefit of design/manufacturing optimizations?

- 50% guardband reduction

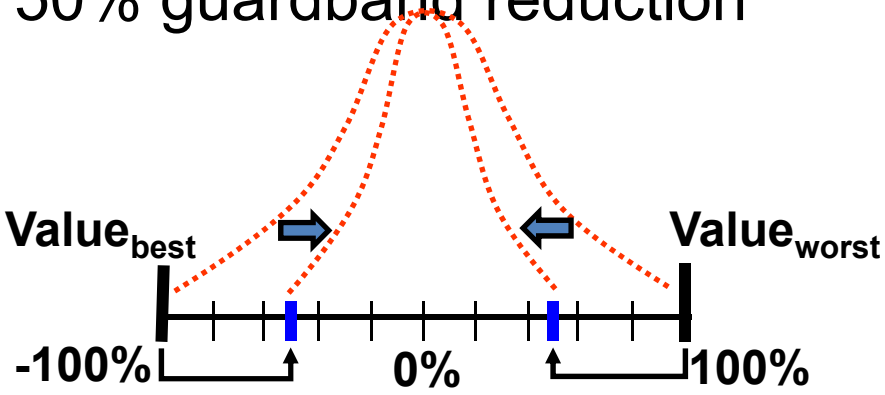

- From delay table analysis:
  - Worst case delay ➔12.5% reduction
- From capacitance table analysis:
  - Worst case cap ➔ 4% reduction

Expected impact of guardband reduction

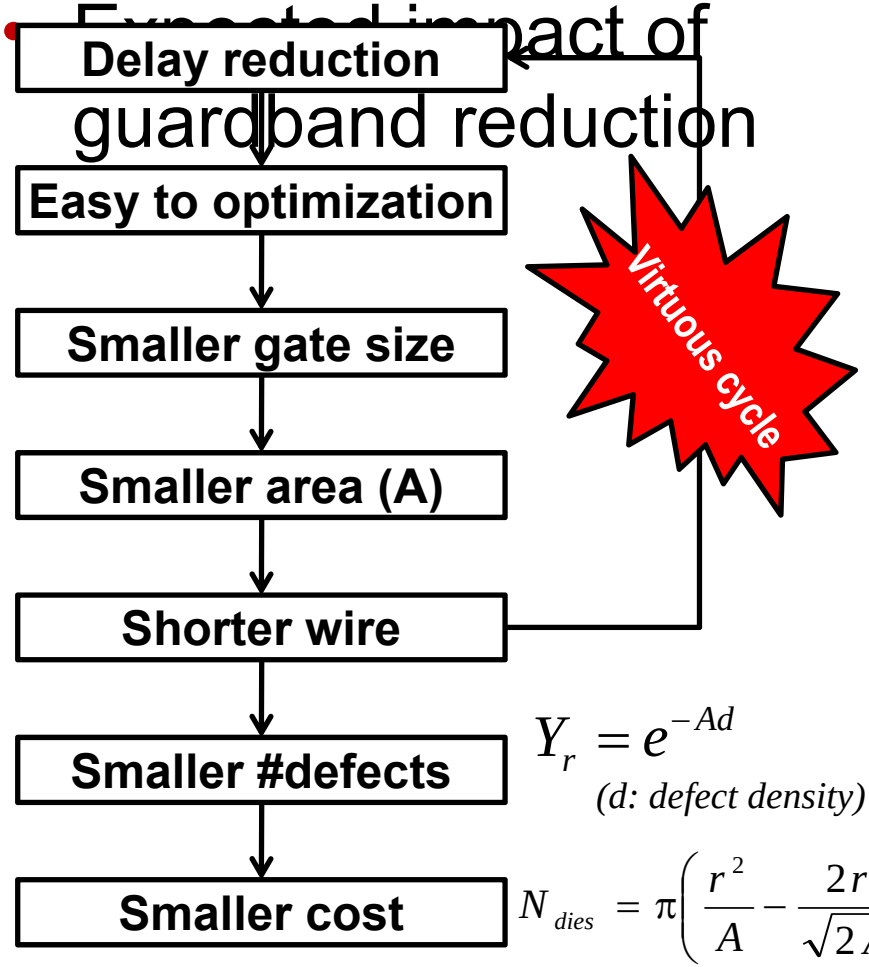

$$Y_r = e^{-Ad}$$

(d: defect density)

$$N_{dies} = \pi\left(\frac{r^2}{A} - \frac{2r}{\sqrt{2A}}\right)$$

(r: wafer radius)

Design Outcomes from Guardband Reduction

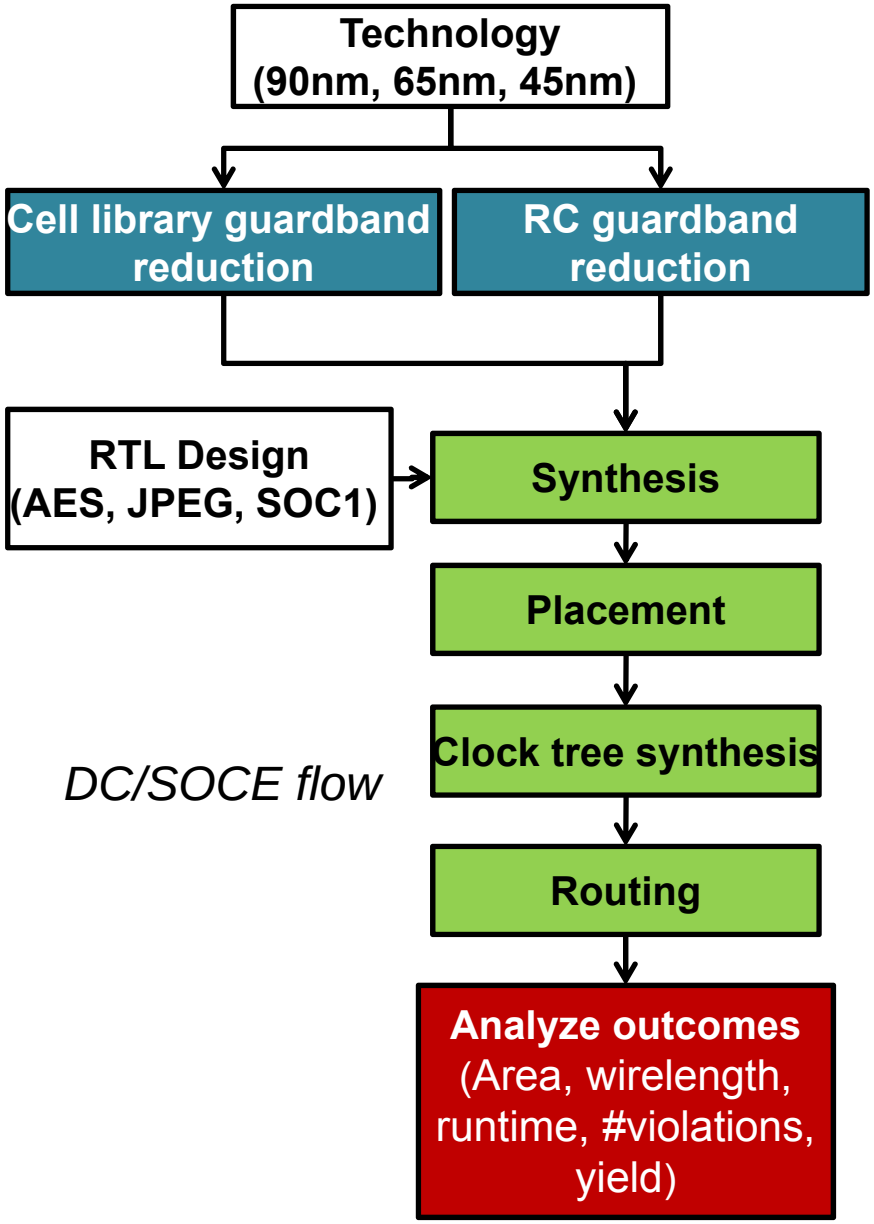

DC/SOCE flow

- 40% guardband reduction
  - Area: 13% reduction
  - Dynamic power: 13% reduction
  - Leakage power: 19% reduction
  - Wirelength: 12% reduction
  - SP&R runtime: 28% reduction
  - #Timing viols.:100% reduction
  - #Good dies (w/o process enhancement): 4% increase

Guardband has very real costs!

(AVS can recover power from overdesign, but not area…)

“Four Horsemen of Dark Silicon Apocalypse”

- Cf. recent talks by Prof. Michael Taylor, UCSD
- **Shrinking Horseman**: “Area is expensive. Chip designers will just build smaller chips instead of having dark silicon in their designs”
- **Dim Horseman**: “We will fill the chip with homogeneous cores that would exceed the power budget but we will underclock them (spatial dimming), or use them all only in bursts (temporal dimming) “
- **Specialized Horseman**: “We will use all of that dark silicon area to build specialized cores, each tuned for the task at hand (10-100x more energy efficient), and only turn on the ones we need…”
- **Deus Ex Machina Horseman**: “MOSFETs are the fundamental problem.”

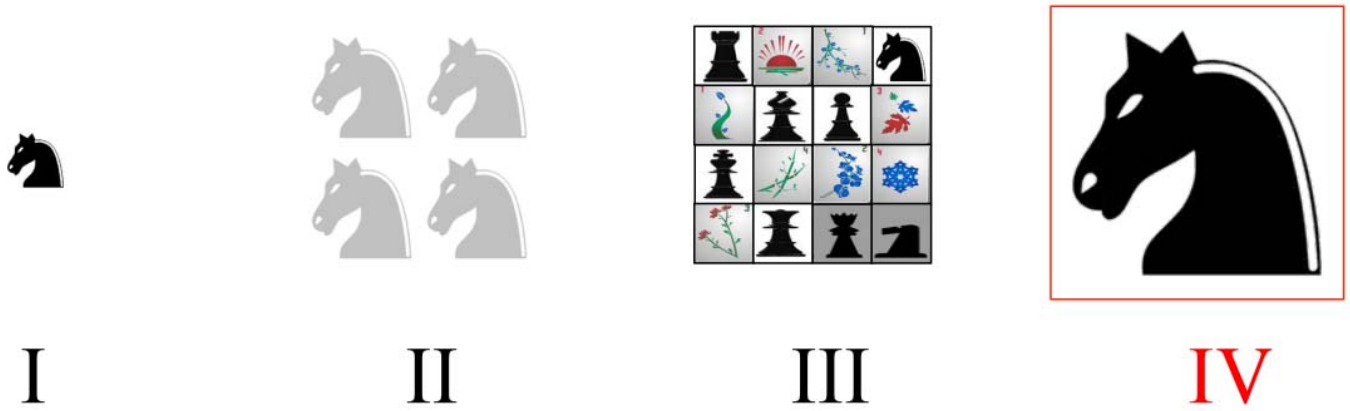

Michael B. Taylor, “Is Dark Silicon Useful? Harnessing the Four Horsemen of the Coming Dark Silicon Apoclypse”, in Proc. DAC, 2012, pp. 1131-1136.

"Dark Silicon" Analysis in 2001 ITRS

- Power management gap ⇒ amount of (switched) logic content in an SOC goes to zero
- Challenge: keeping the chip value above zero