VLSI-2:
Variability Aware Circuit Design

Alan Drake
Austin Design Center
TSMC

Agenda

- Why do we care about variability
- What are the sources of variability
  - Process
  - Temperature
  - Voltage
  - Noise
  - Ageing
- What do we do about it
  - Margining
  - Guardband
  - Critical Path Monitor Examples

Motivation

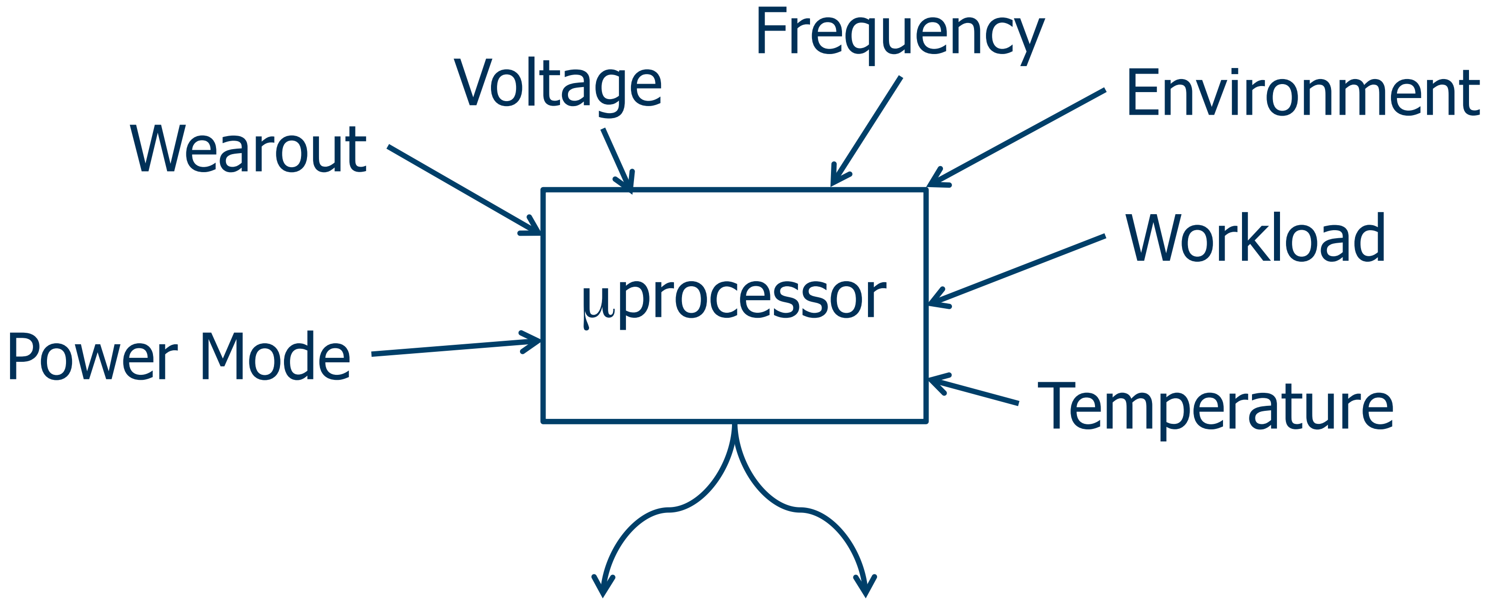

Frequency

Voltage

Environment

Wearout

µprocessor

Workload

Power Mode

Temperature

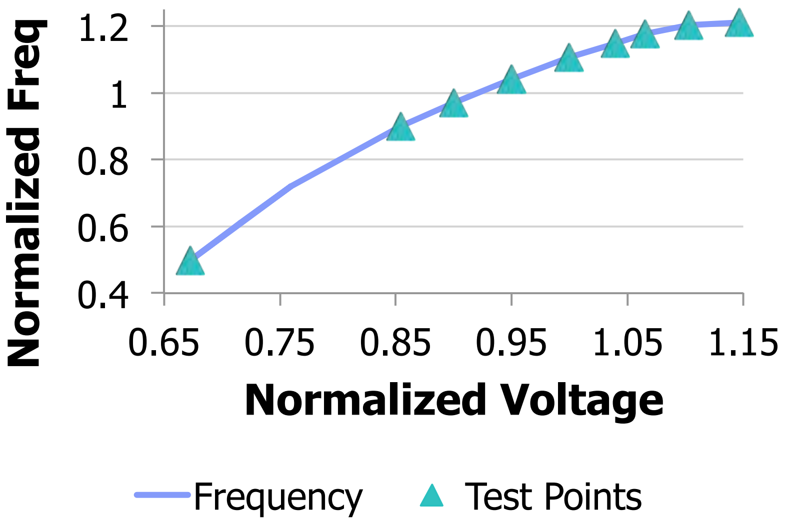

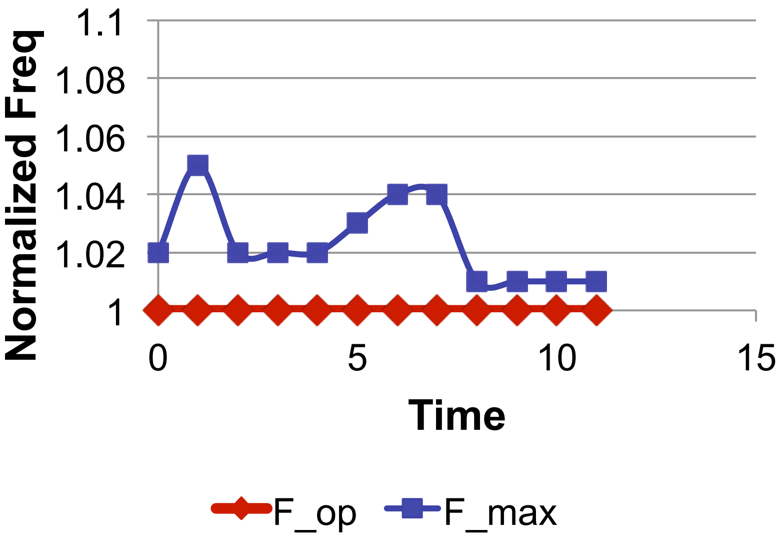

Process Variation

- Fixed noise
- Two types: die-to-die and within-die
- Variation seems to be increasing with scaling
- Typically measured statically from initial technology development through production
- Sensors used for measurement
  - Ring Oscillators
  - Characterization Arrays

Ring Oscillators

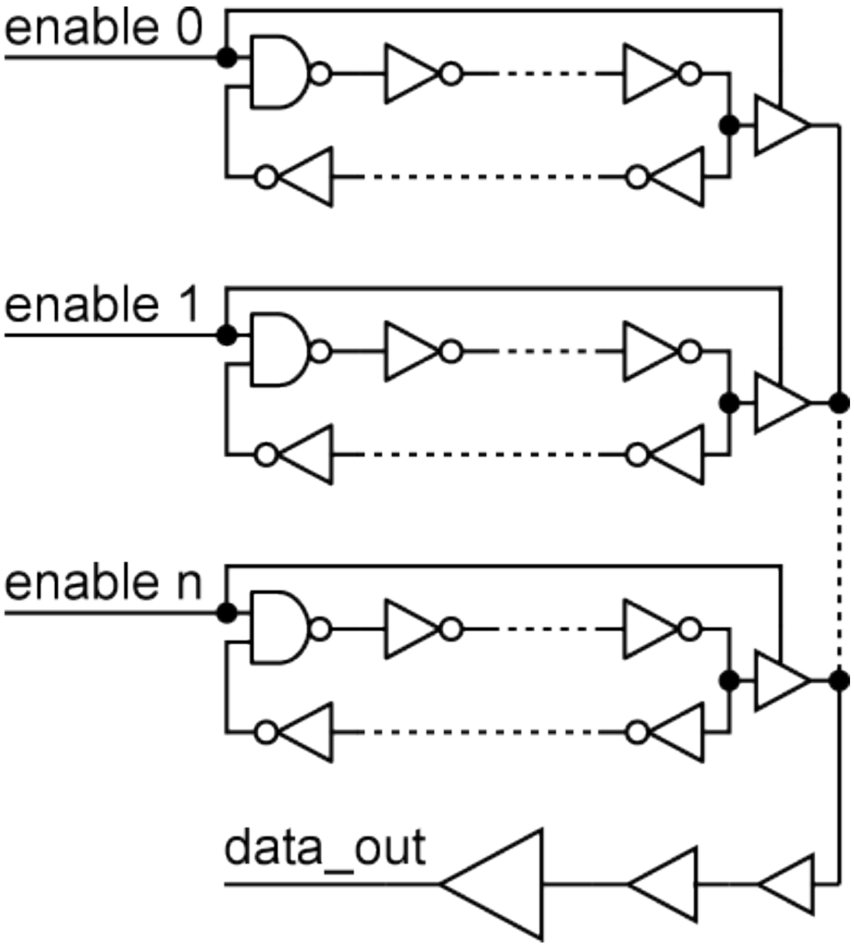

- Provide average value of random noise
- Used to accumulate non-random variation to a measureable quantity
- Noise to frequency converter

Processor Performance Variability

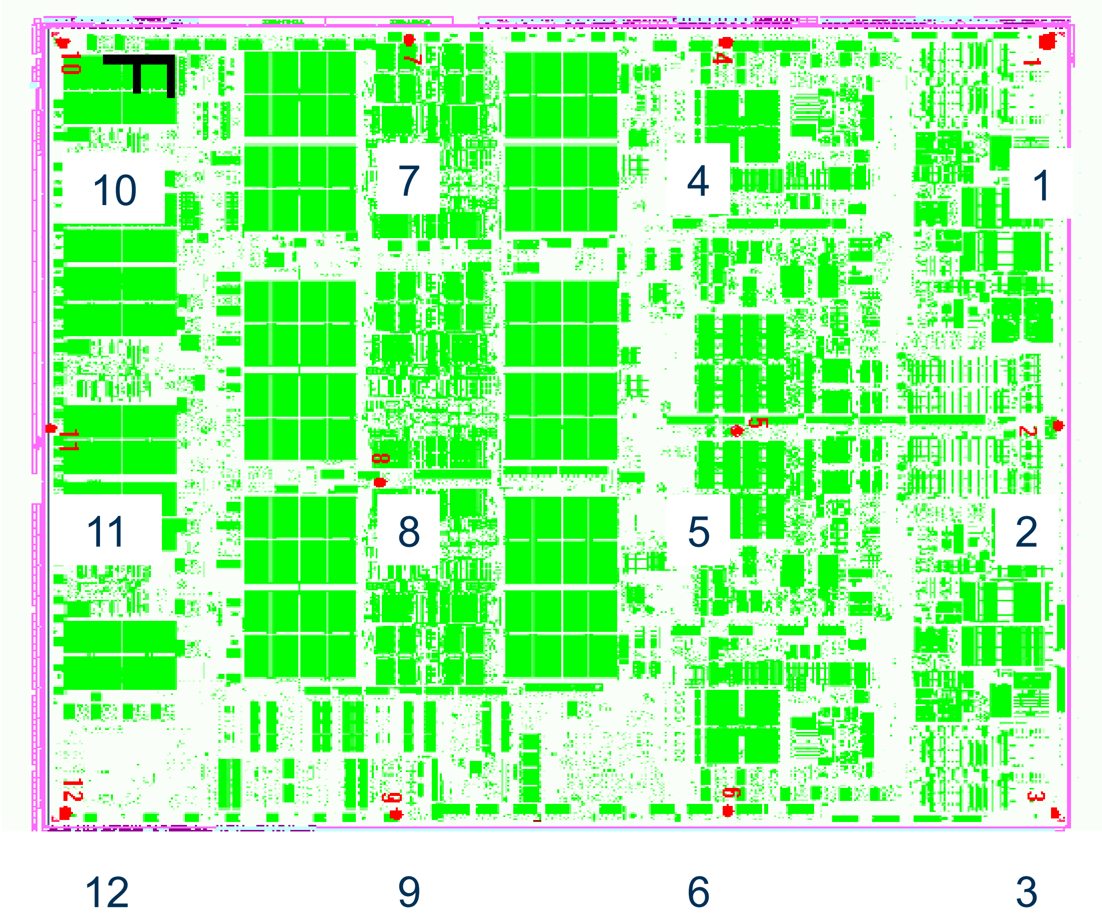

Chip map with Ring Oscillator locations

- Chip has 12 ring oscillators distributed across the die, and individually measurable.

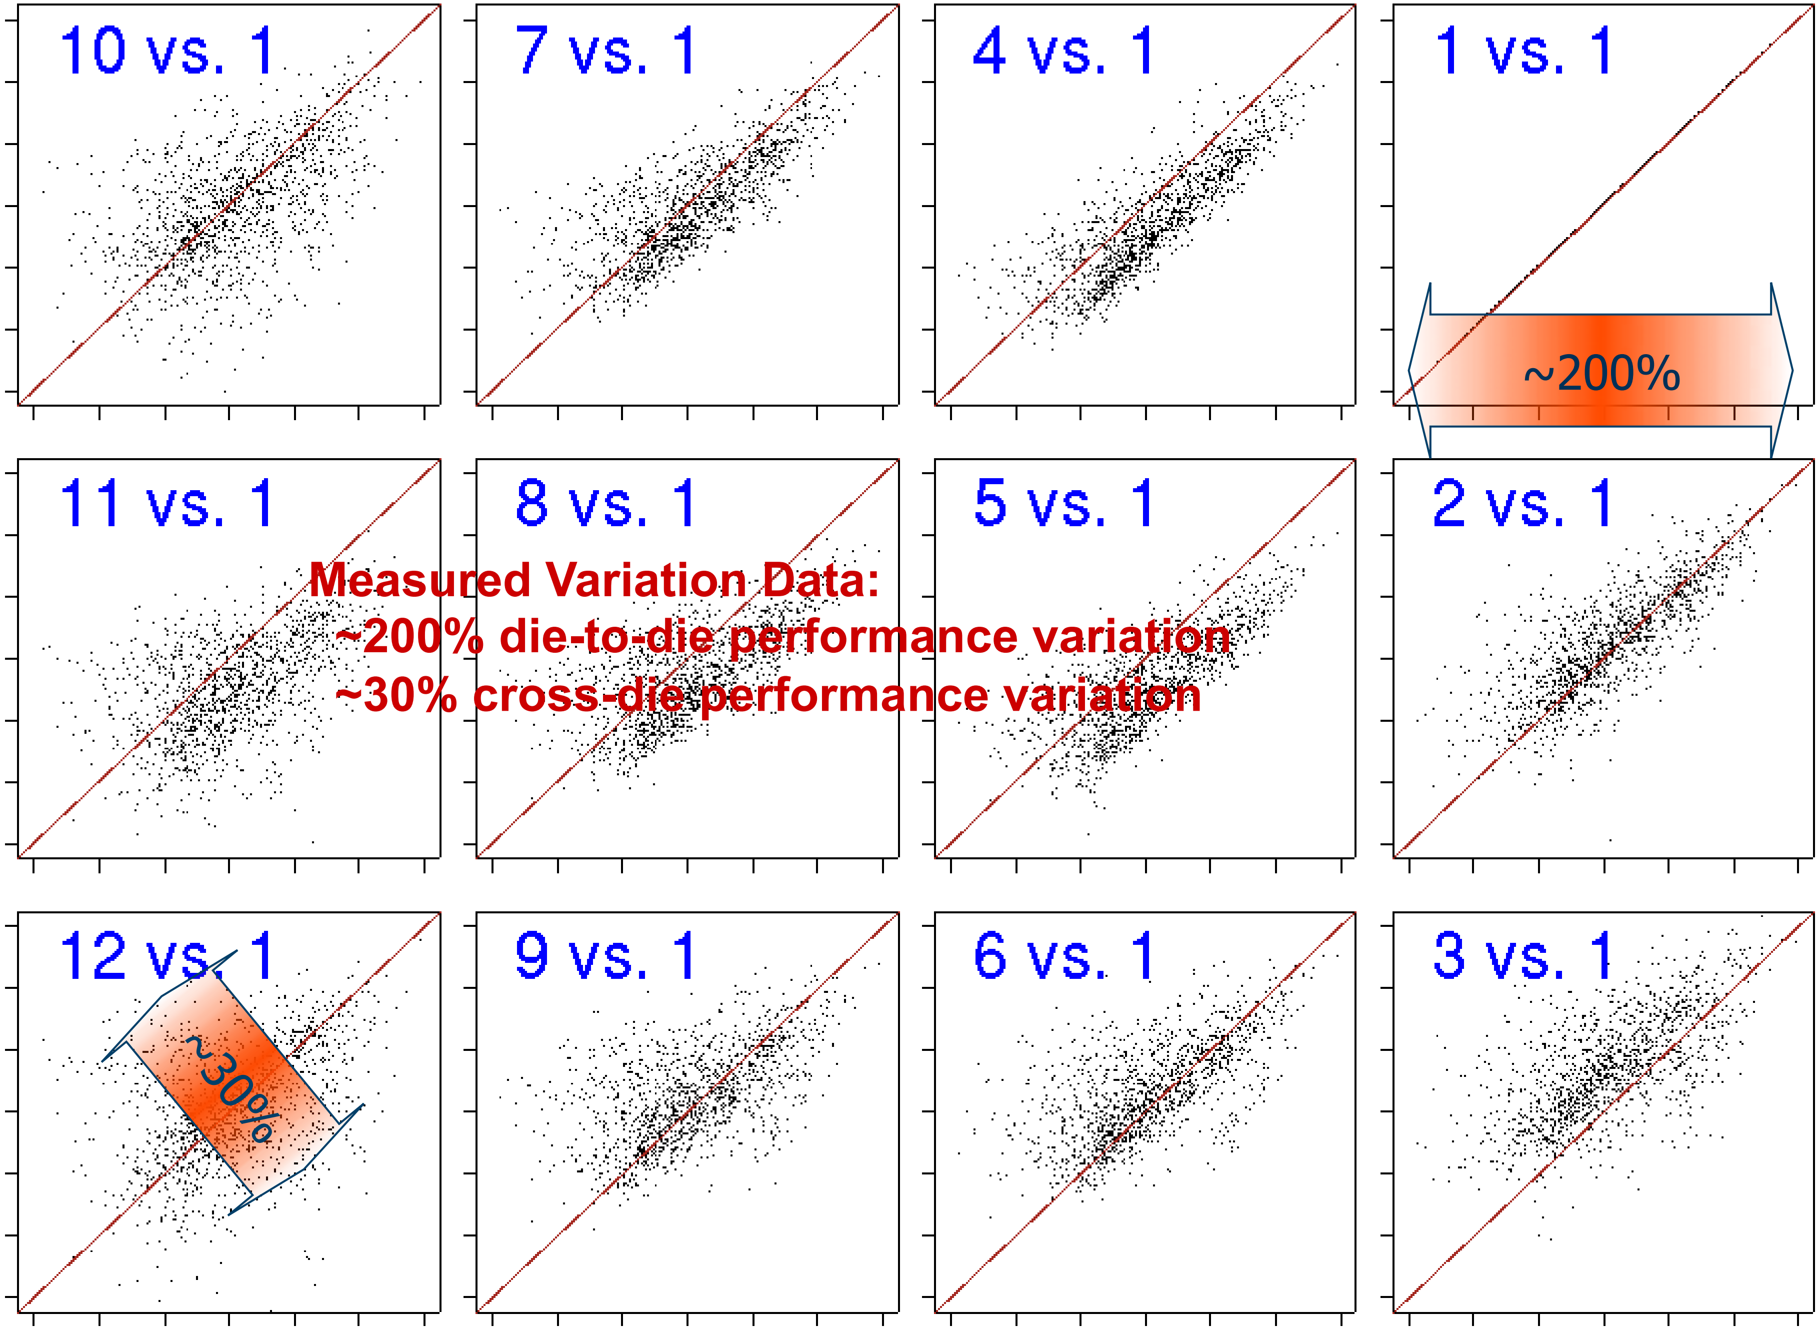

Measured Variation Data:
~200% die-to-die performance variation
~30% cross-die performance variation

V_T Characterization Array

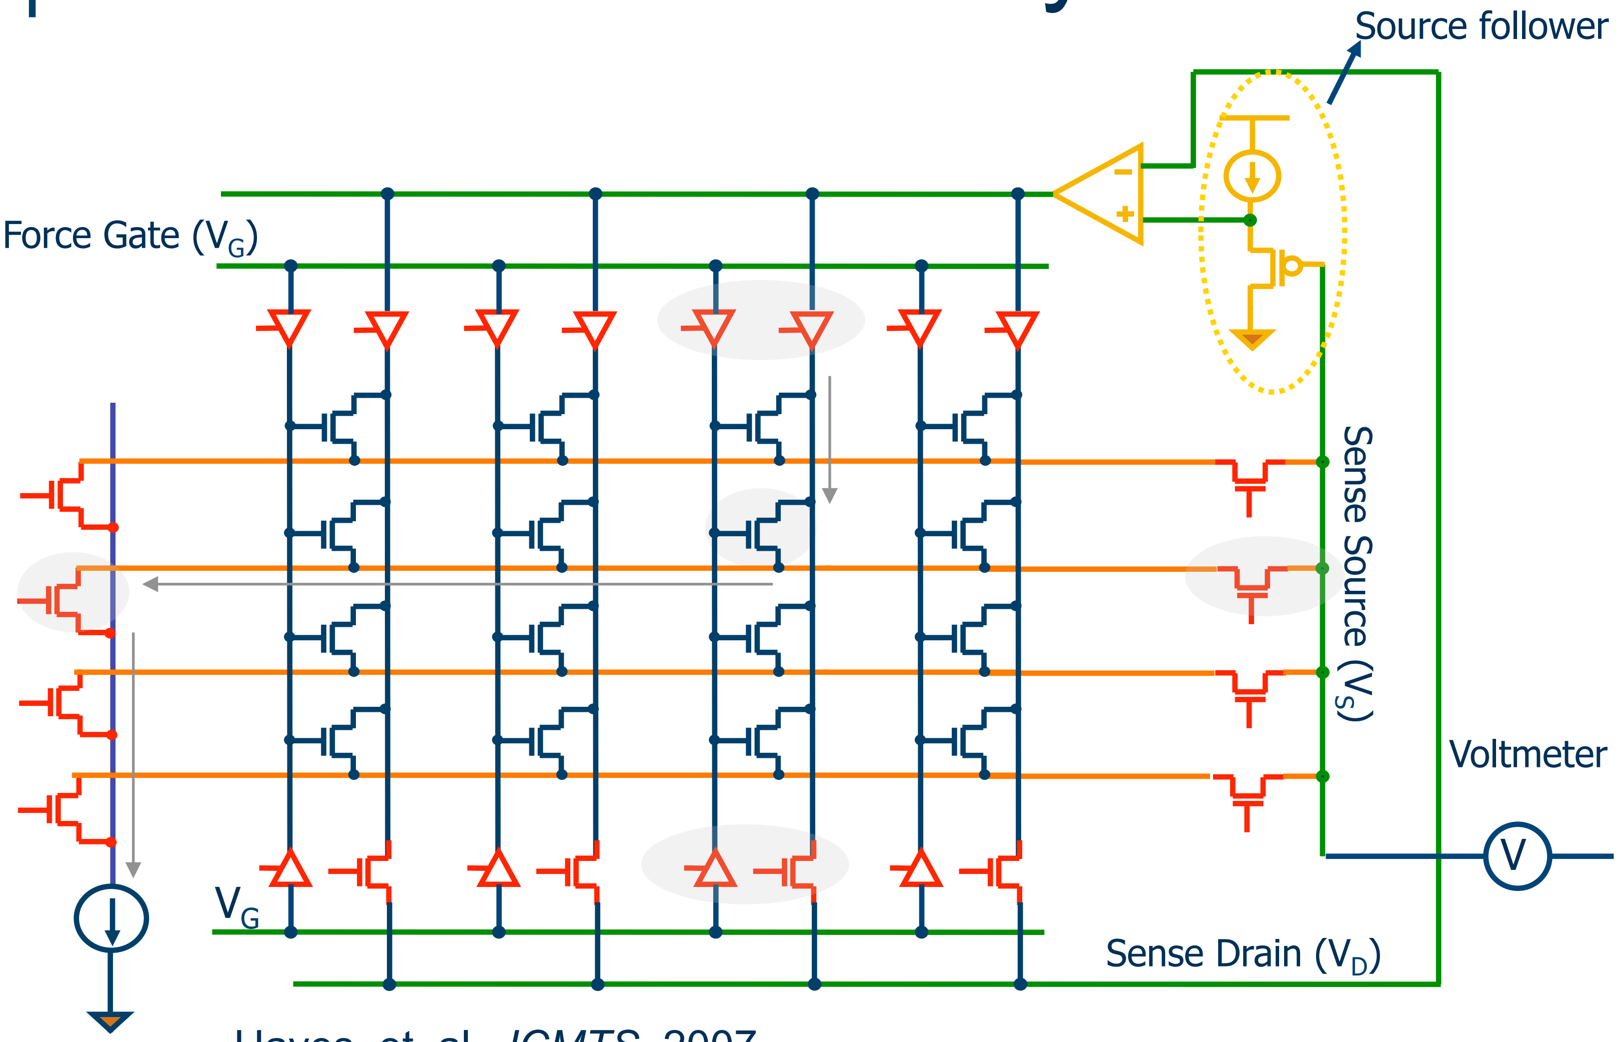

Hayes, et. al., *ICMTS*, 2007

Measured V_T Variation in 65 nm PD-SOI

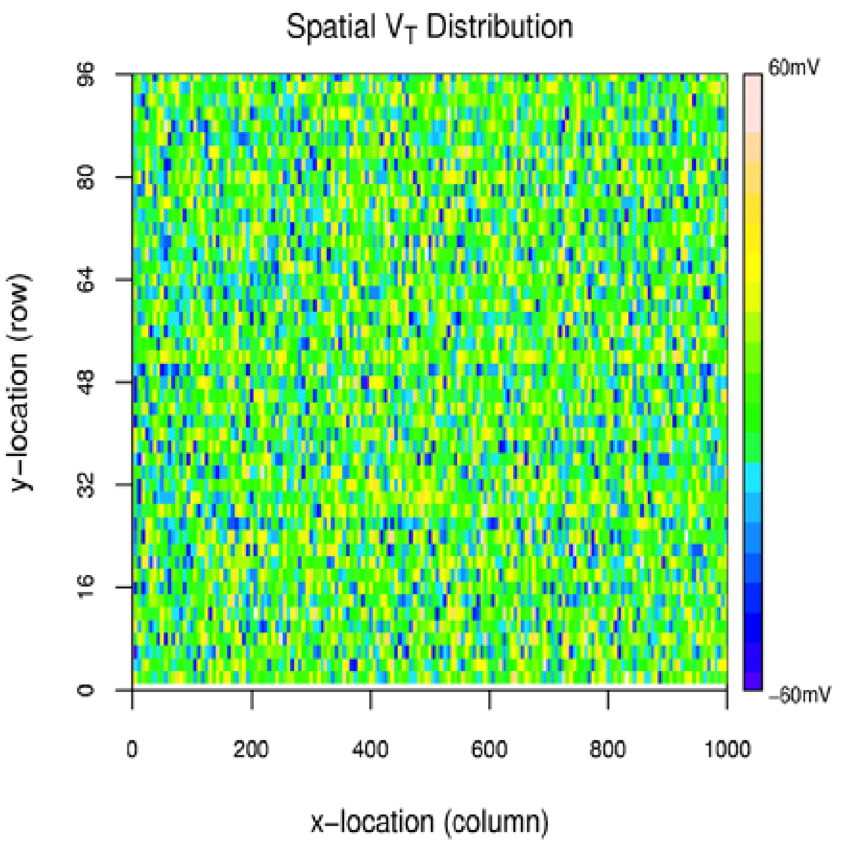

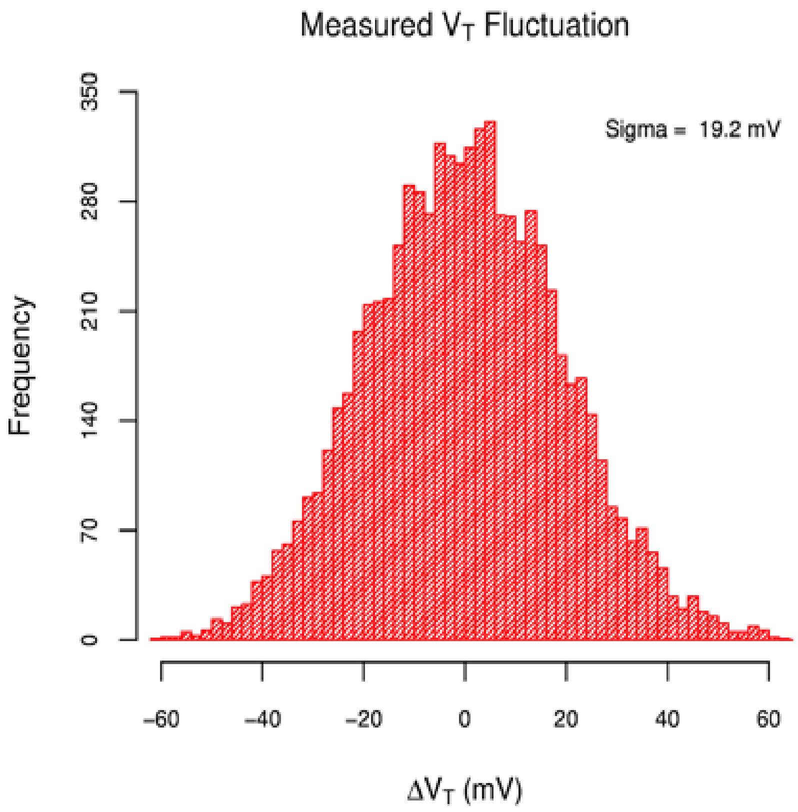

- 96,000 DUTs
- No Spatial correlation across the array

Hayes, et. al., *ICMTS*, 2007

- 8000 DUTs
- Random variability across the array

Temperature

- Differences larger than 50°C possible
- Results from static power, dynamic power, and external environment
- Slow moving signal
- Constant over a small distance due to heat spreading
  - SOI and FINFETs complicate this due to thermal resistance of oxide layer
- Accuracy ranges of on-chip sensors range from about 5°C to less than 0.1°C
- Sensors often used for measuring temperature
  - PTAT
  - Resistance
  - Delay line (covered in a later section)

Device Power Variability

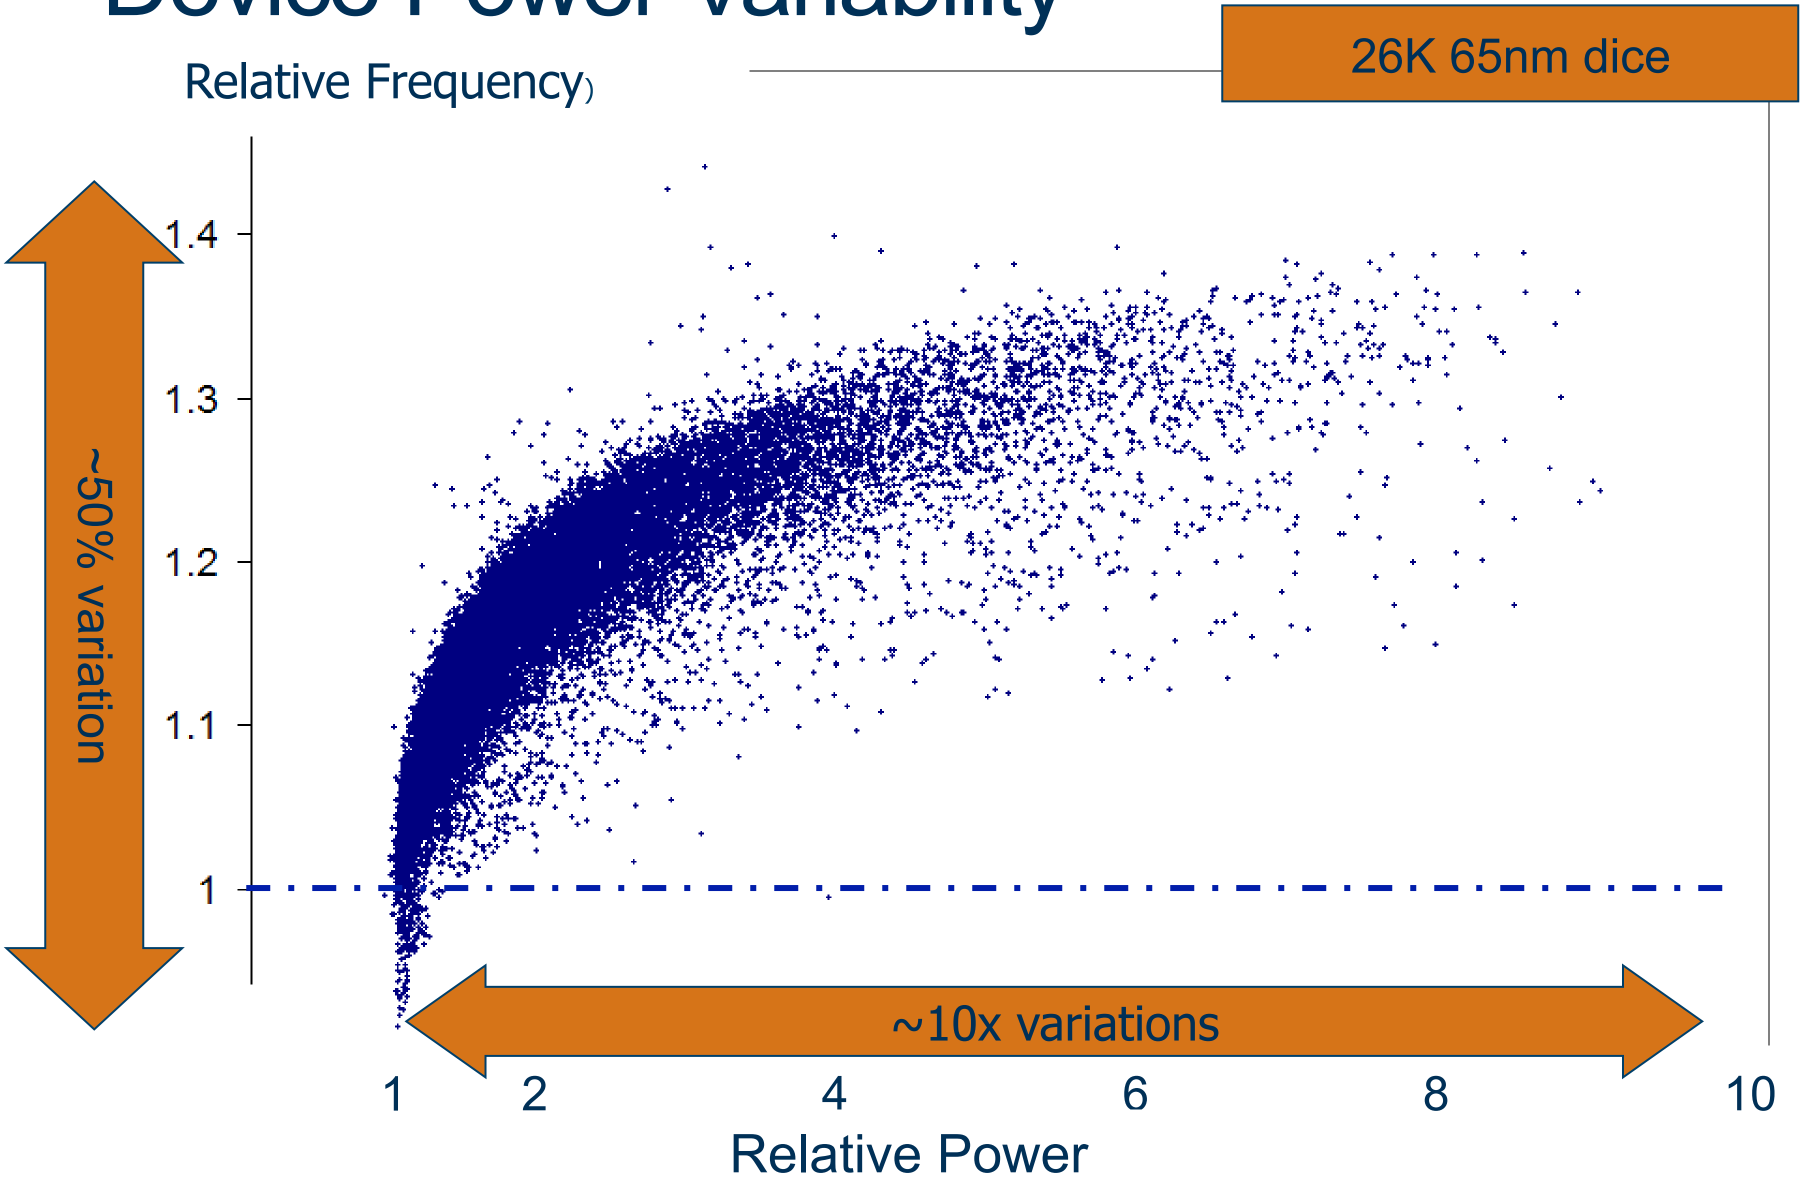

Voltage Droop - Temperature Map

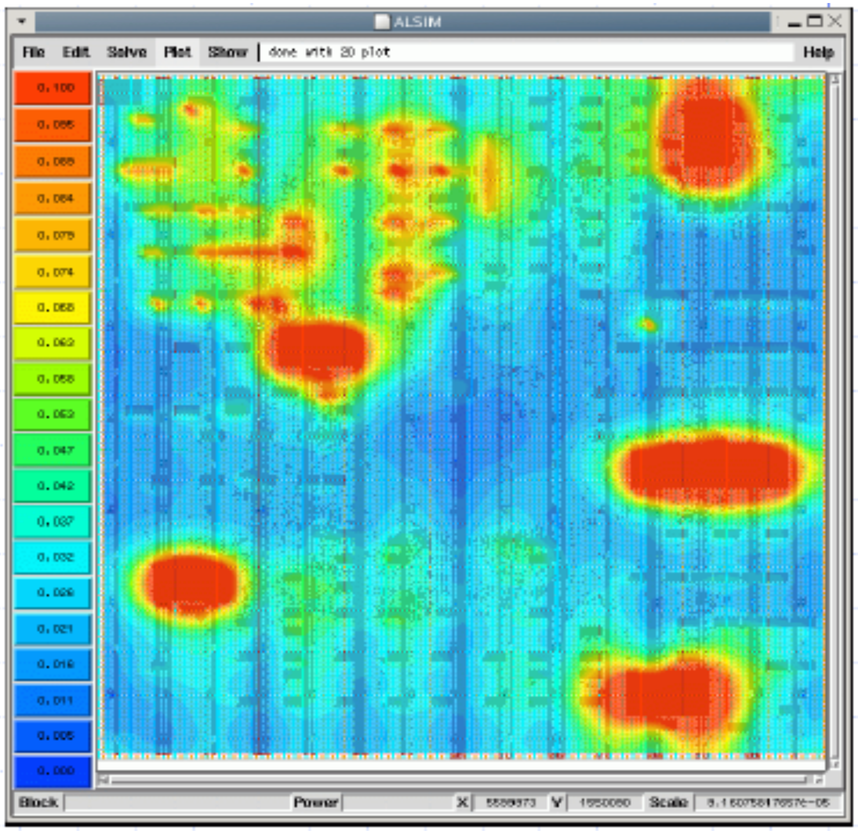

done with 2D plot

Power supply droop map

Source: Sani Nassif, IBM

- **Supply droop caused by high load currents, indicating areas of high temperature.**

Temperature Dependence of Circuit Elements - Band gap

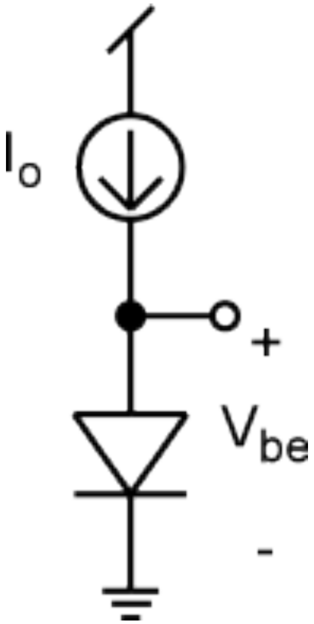

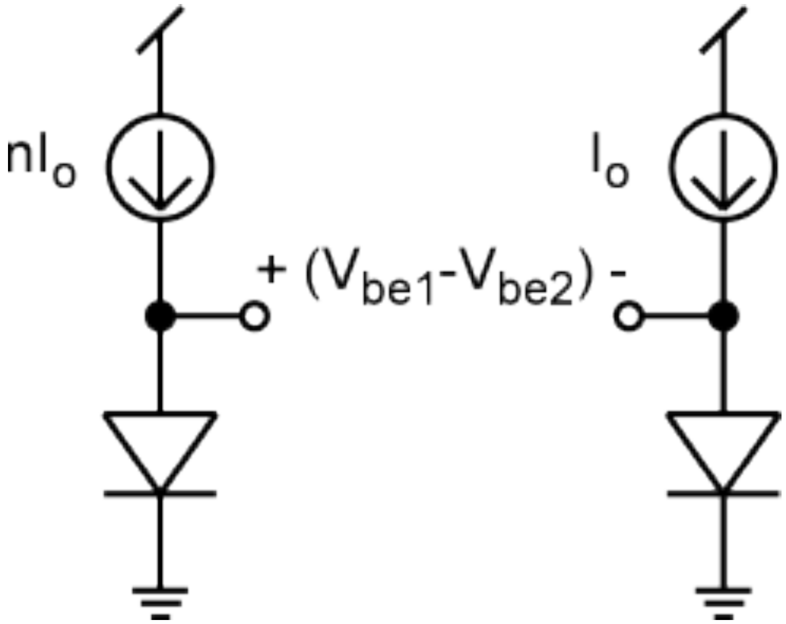

-2 mV/°C

$$\frac{k}{q}\ln n$$

Temperature Dependence of Circuit Elements - Resistance

$$R = R_0\left[1 + \alpha\left(T - T_0\right)\right]$$

Gold	3.4×10^{-3}/°C
Copper	3.9×10^{-3}/°C
Aluminum	3.9×10^{-3}/°C
Tungsten	4.5×10^{-3}/°C

Temperature Dependence of Circuit Elements – Gate Delay

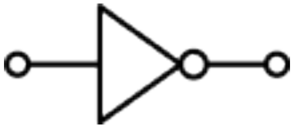

- Temperature affects carrier mobility
- Long-channel approximation
- First order delay model

$$\beta \propto T^{-1.5}$$

$$I_D \approx \alpha T^{-1.5}$$

$$t = \frac{CV_{DD}}{I_D}$$

$$\frac{\partial t}{\partial T} = C\frac{\left(\frac{\partial V_{DD}}{\partial T} I_D - V_{DD}\frac{\partial I_D}{\partial T}\right)}{I_D{}^2}$$

$$\frac{\partial t}{\partial T} = \frac{1.5CV_{DD}}{\alpha}\sqrt{T}$$

Sample PTAT Current Reference

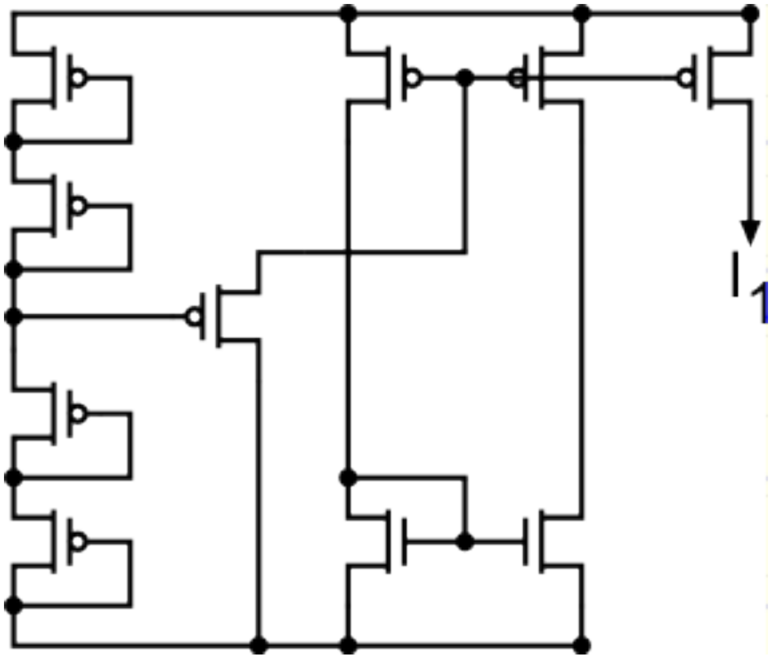

- 65 nm PD-SOI
- Temperature sensitivity is 0.133 μA/°C
- Voltage sensitivity is 0.05 μA/V
- μ = 28.0 μA
- σ = 2.0 μA

Fadi Gebara, IBM Austin Research Lab, fhgebara@us.ibm.com

Sample PTAT Current vs. T

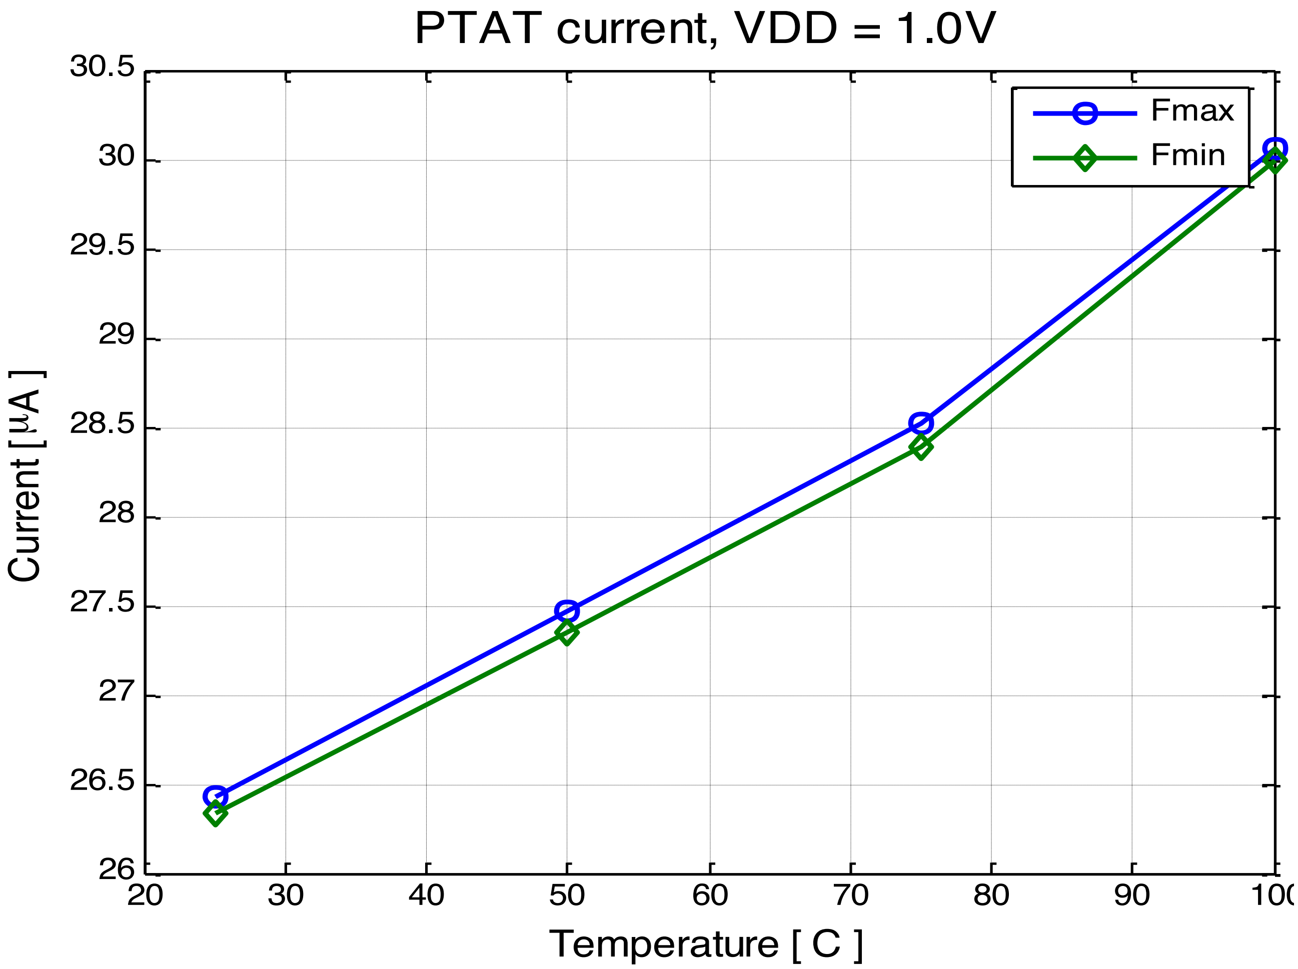

Fadi Gebara, IBM Austin Research Lab, fhgebara@us.ibm.com

PTAT Based Temperature Sensor

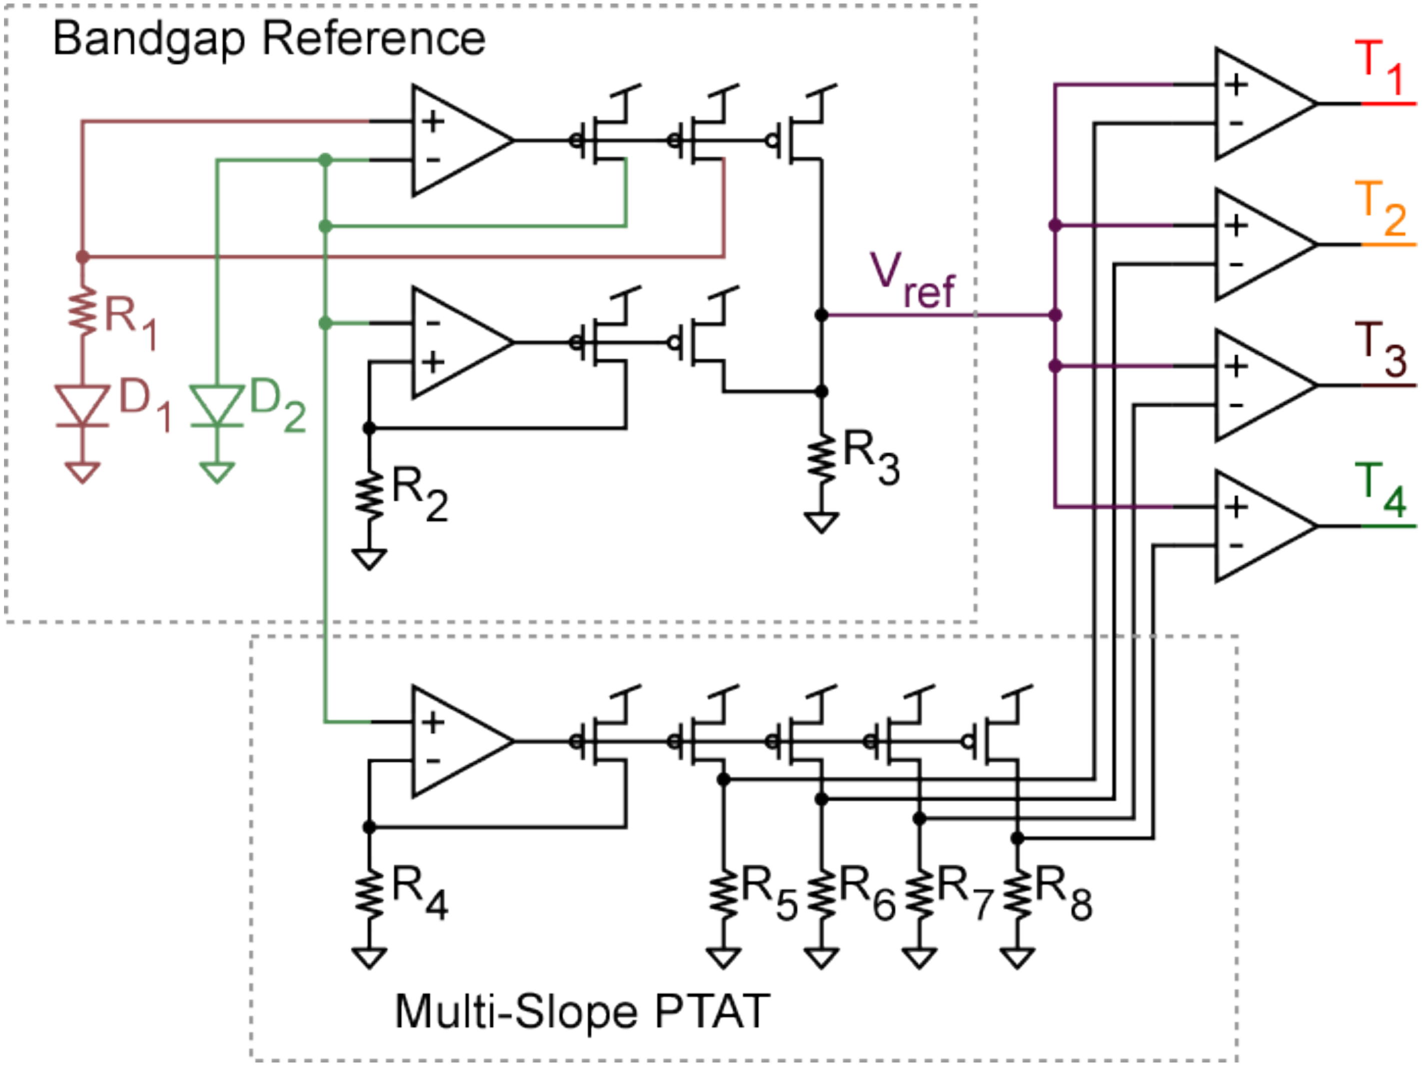

Hsu, et. al., US Patent 6,876,250 B2, 2005

Temperature Sensor Data Conversion

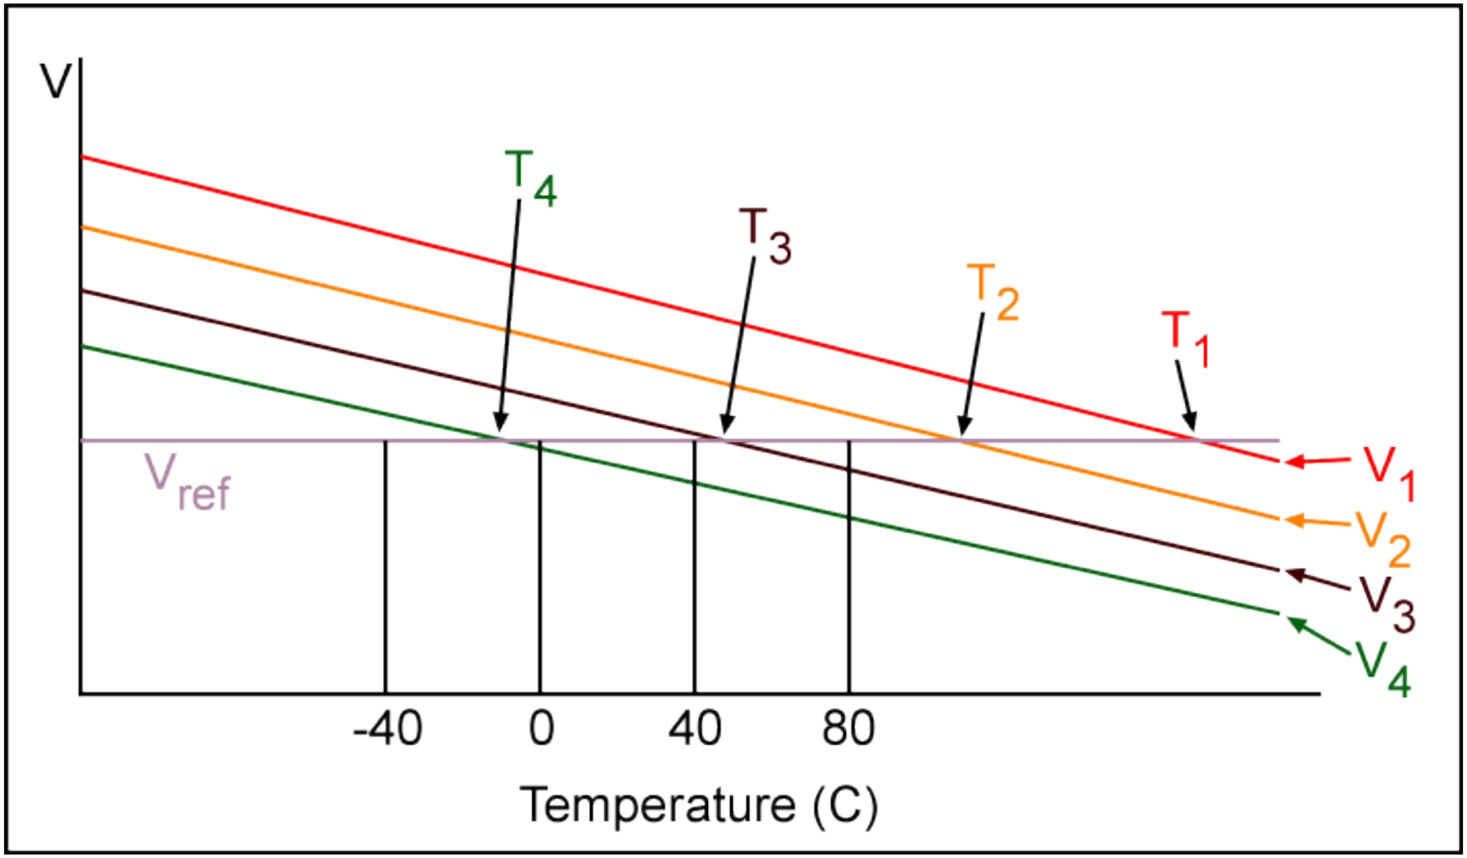

Hsu, et. al., US Patent 6,876,250 B2, 2005

Voltage

- Fast transient noise
- Highly dependent on workload in modern microprocessors
- Colored by clock frequency and repetitive data patterns
- Manifests in timing uncertainty and clock jitter
- Quickly propagates across power grid
- Covered sensors
  - Flash A/D converter
  - Delay line

V_{DD} Noise Coupling

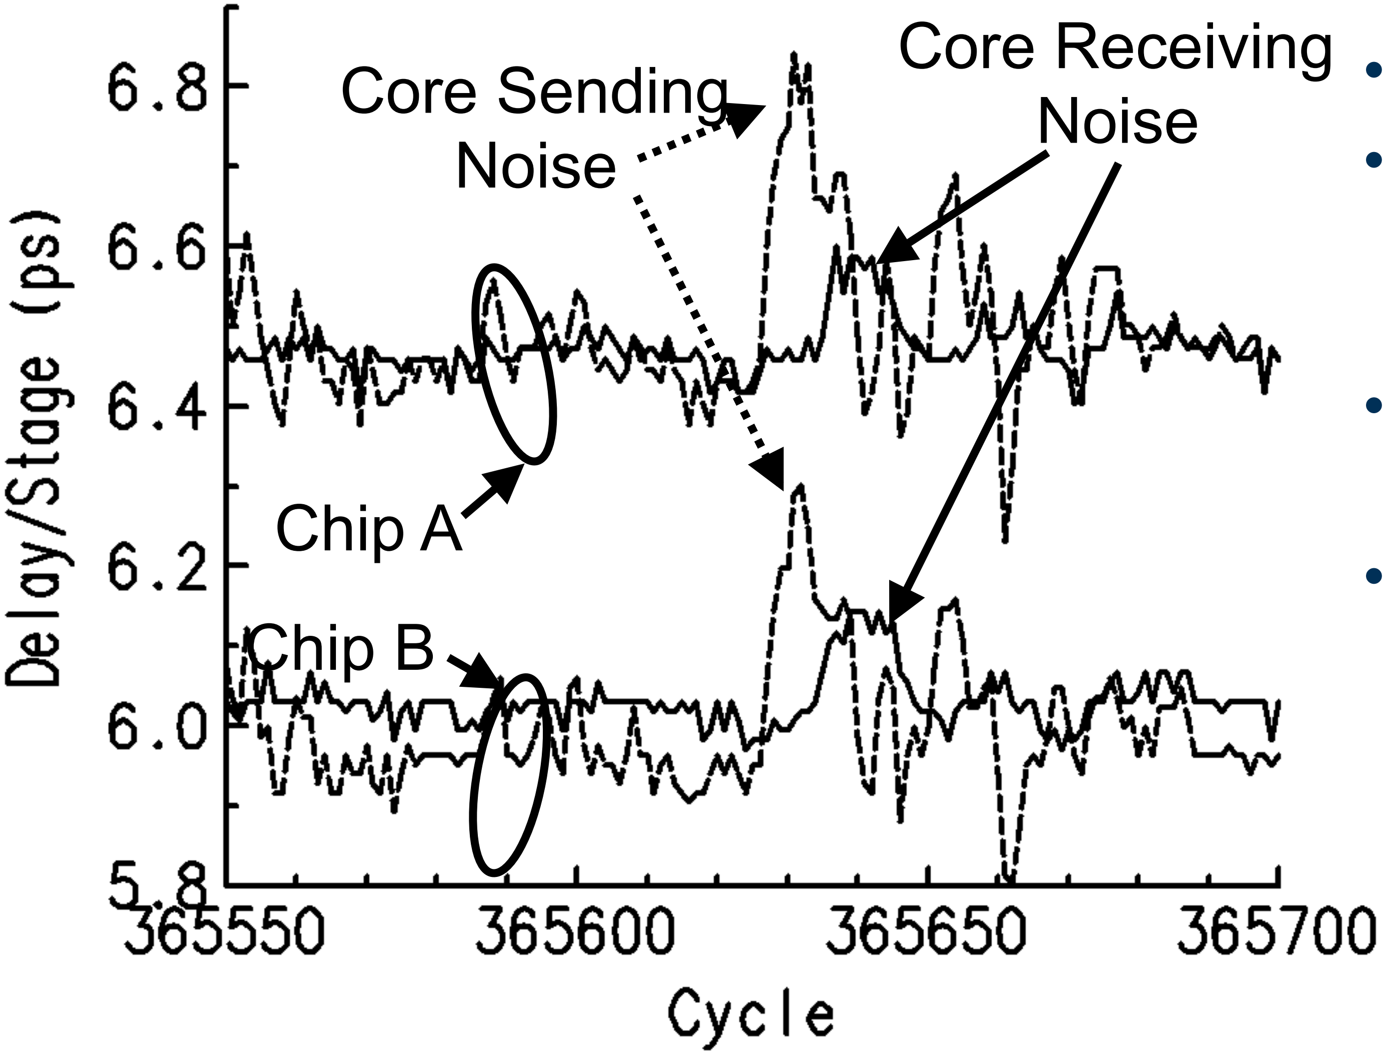

Core Sending Noise

Core Receiving Noise

Chip A

Chip B

- 2 chips
- Sending core is running code
- Receiving core is idle
- Noise is attenuated at receiving core

Norman James et al, ISSCC 2007

Noise Traveling Across Chip

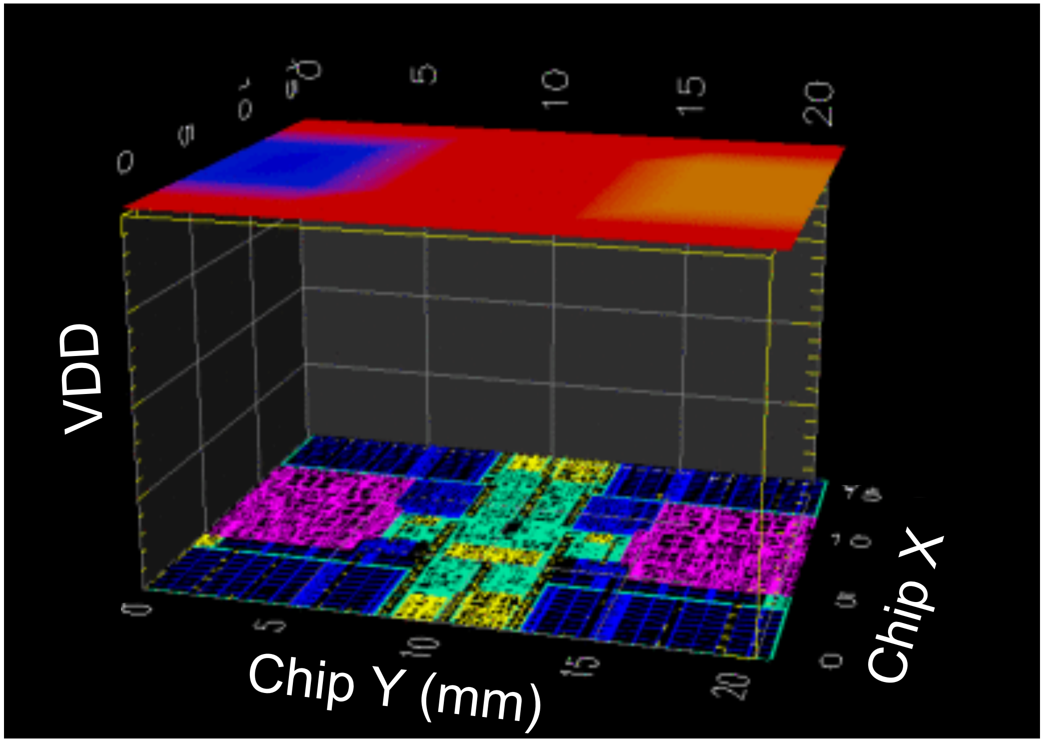

- Noise travels from active core to idle core
- ~4ns delay between cores
- Noise is attenuated at quiet core

Thomas Strach, IBM Systems and Technology Group, Boeblingen, Germany

Phillip Restle, T. J. Watson Research Lab, Yorktown Heights, NY

Worst Case – “The Perfect Storm”

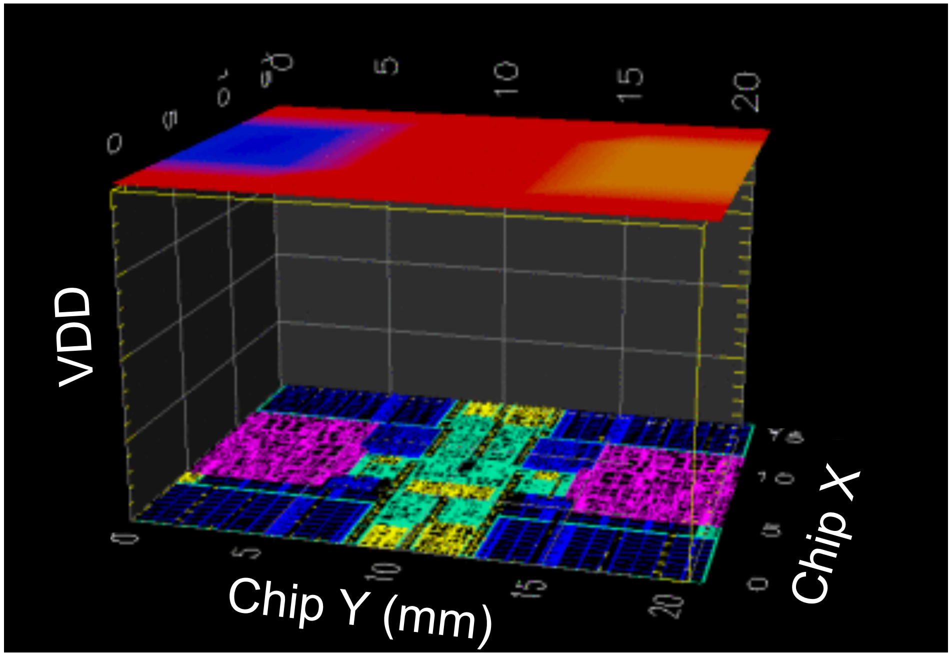

- Same code running on both cores
- Right core delayed by 4ns
- Noises combine in right core to give worst case

Thomas Strach, IBM Systems and Technology Group, Boeblingen, Germany

Phillip Restle, T. J. Watson Research Lab, Yorktown Heights, NY

Skitter (Skew and Jitter)

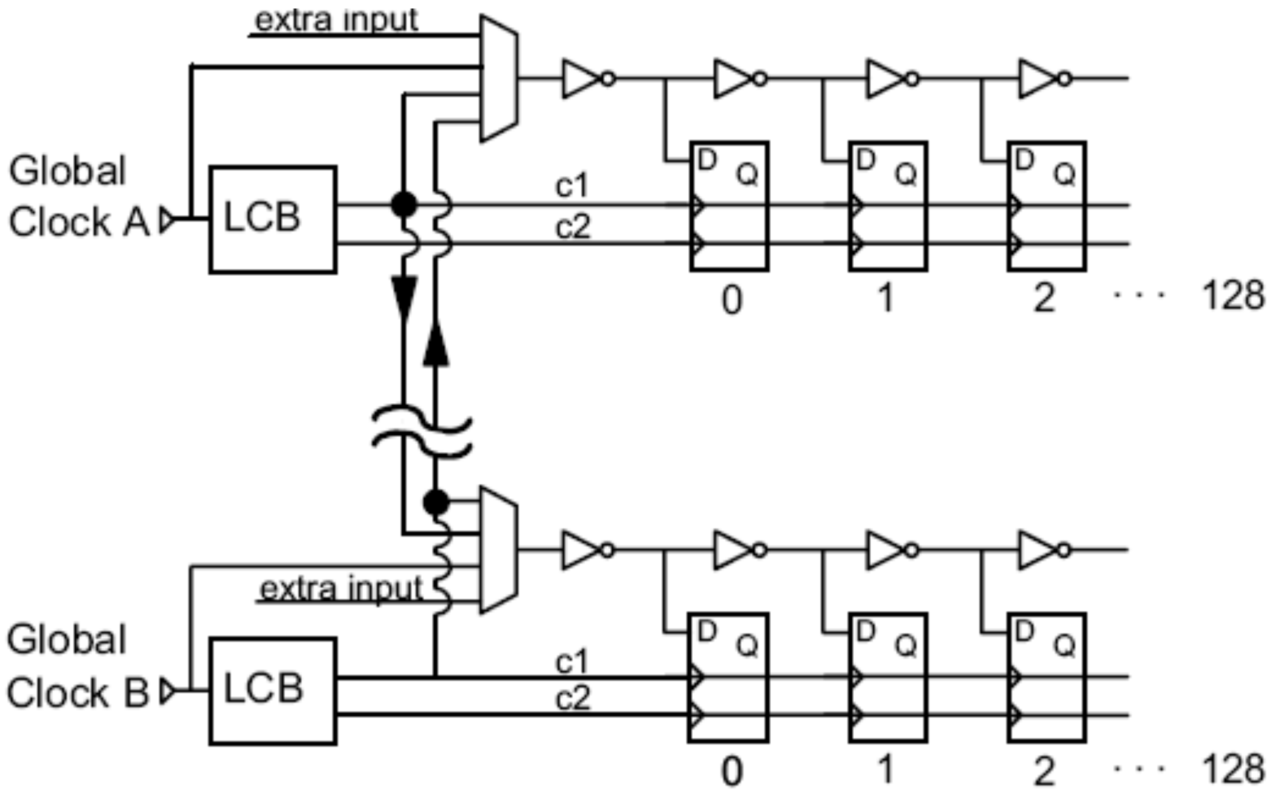

Restle, et. al., *ISSCC*, 2004

V_{DD} Noise Flash Converter

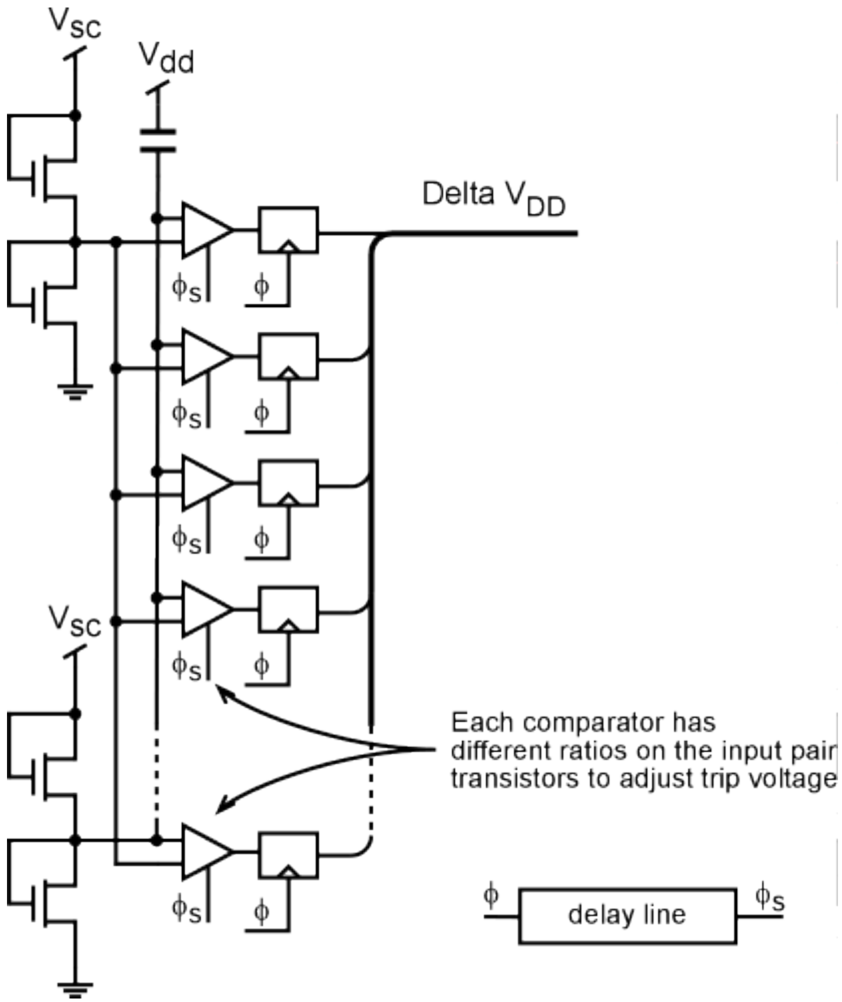

Each comparator has different ratios on the input pair transistors to adjust trip voltage

- Couples noise from V_{dd}
- Only one reference voltage saves power
- Reference divides V_{sc} by two

Negative Bias Temperature Instability

- Caused by large negative gate-source voltage in PMOS
- Transient phenomenon that gradually degrades devices
- An important part of device ageing
- There is also an Positive Bias Temperature Instability affect in NMOS devices

Negative Bias Temperature Instability

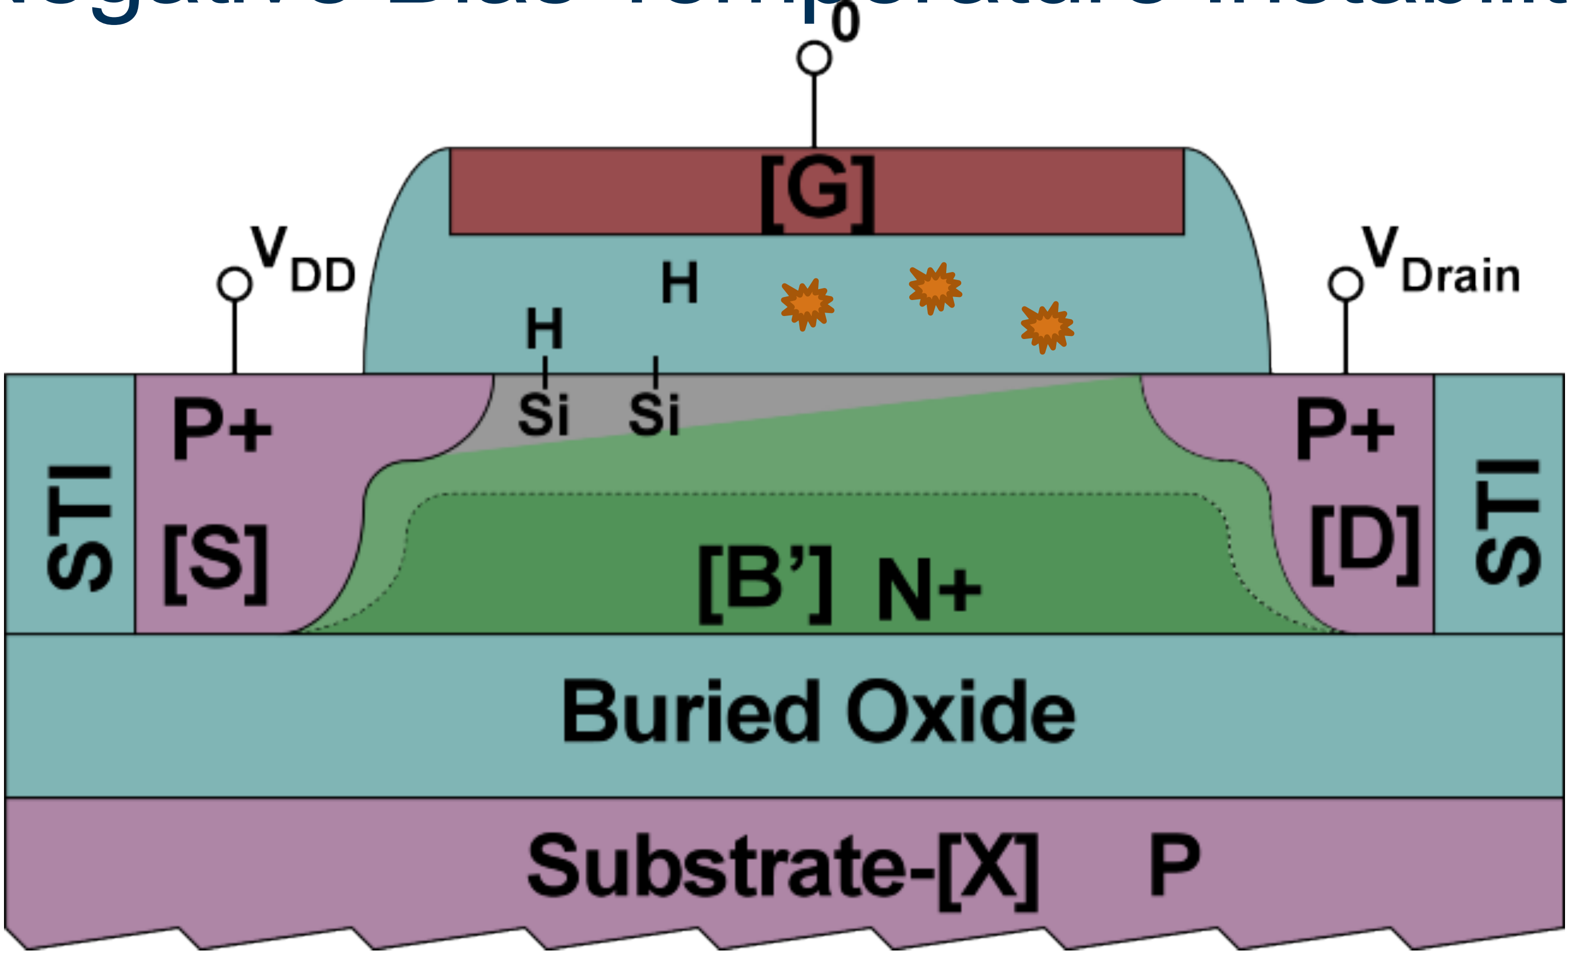

Thought to be created by interface traps (non-recoverable) and internal traps (recoverable).

NBTI Sensor

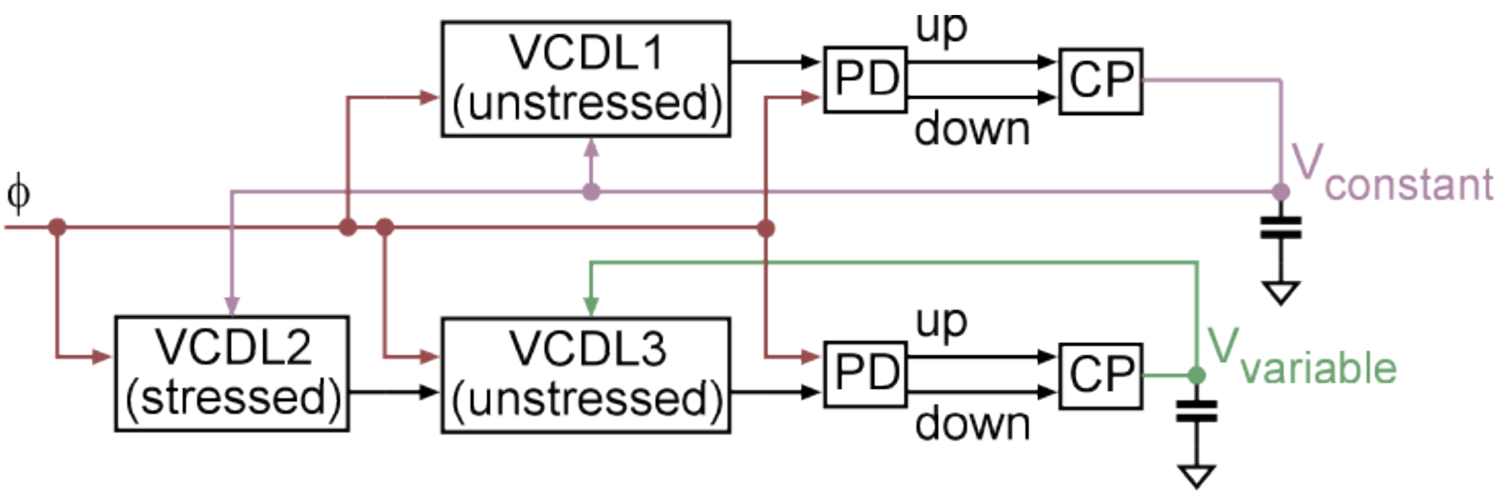

Sensor Block Diagram

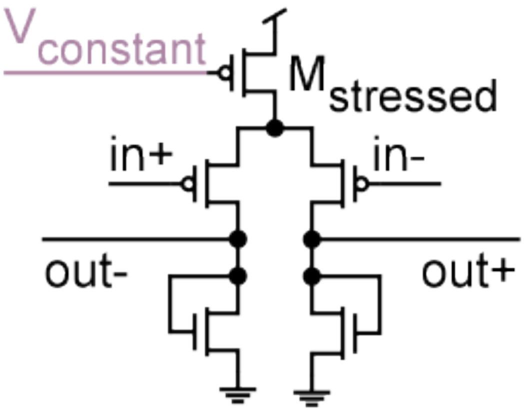

NBTI Stressed Delay Element

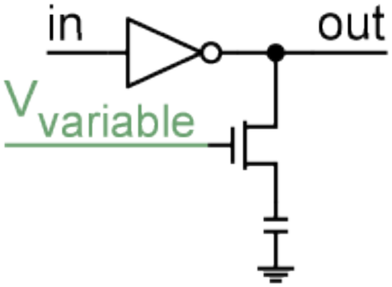

Unstressed Delay Element

Keane et. al., *ISLPED*, 2007

Agenda

- Why do we care about variability
- What are the sources of variability
  - Process
  - Temperature
  - Voltage
  - Noise
  - Ageing
- What do we do about it
  - Margining
  - Guardband
  - Critical Path Monitor Examples

Timing Variation

- Caused by all noise processes previously discussed
- Manifests in clock jitter and skew and path delay changes
- Timing variability only a problem in critical paths

$$T_d + m\sigma(\delta T_d + \delta T_{\theta 1} + \delta T_{\theta 2}) < T_\theta - \delta T_\theta - T_{setup}$$

- Multi-core microprocessors combined make binning more difficult and increase margins
- Covered sensors
  - Delay line

Guard banded voltage/frequency limits

- **Conventional guardband**
  - Static, conservative voltage/frequency margins for potential worst-case conditions
  - Causes unnecessary loss of energy efficiency/ performance during typical server usage

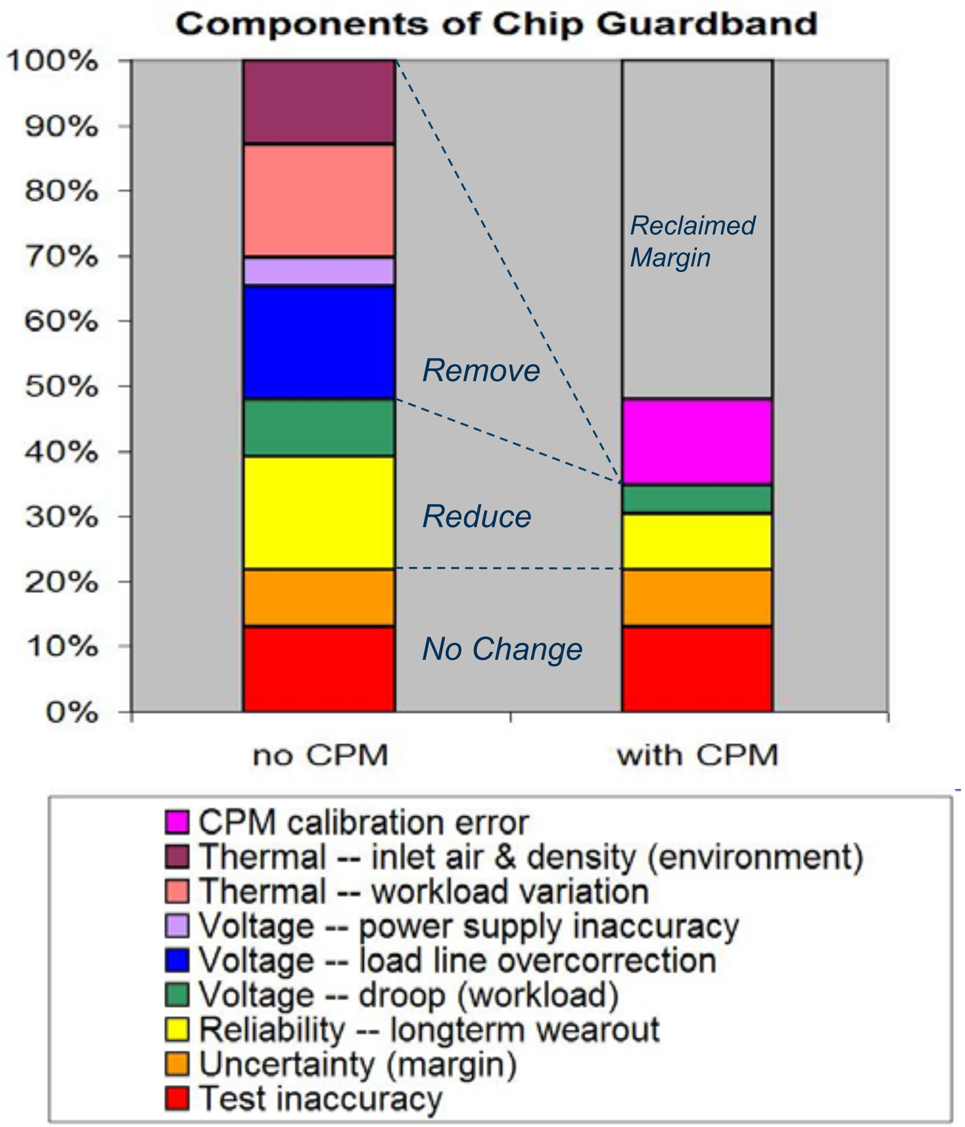

Critical Path Timing Distribution

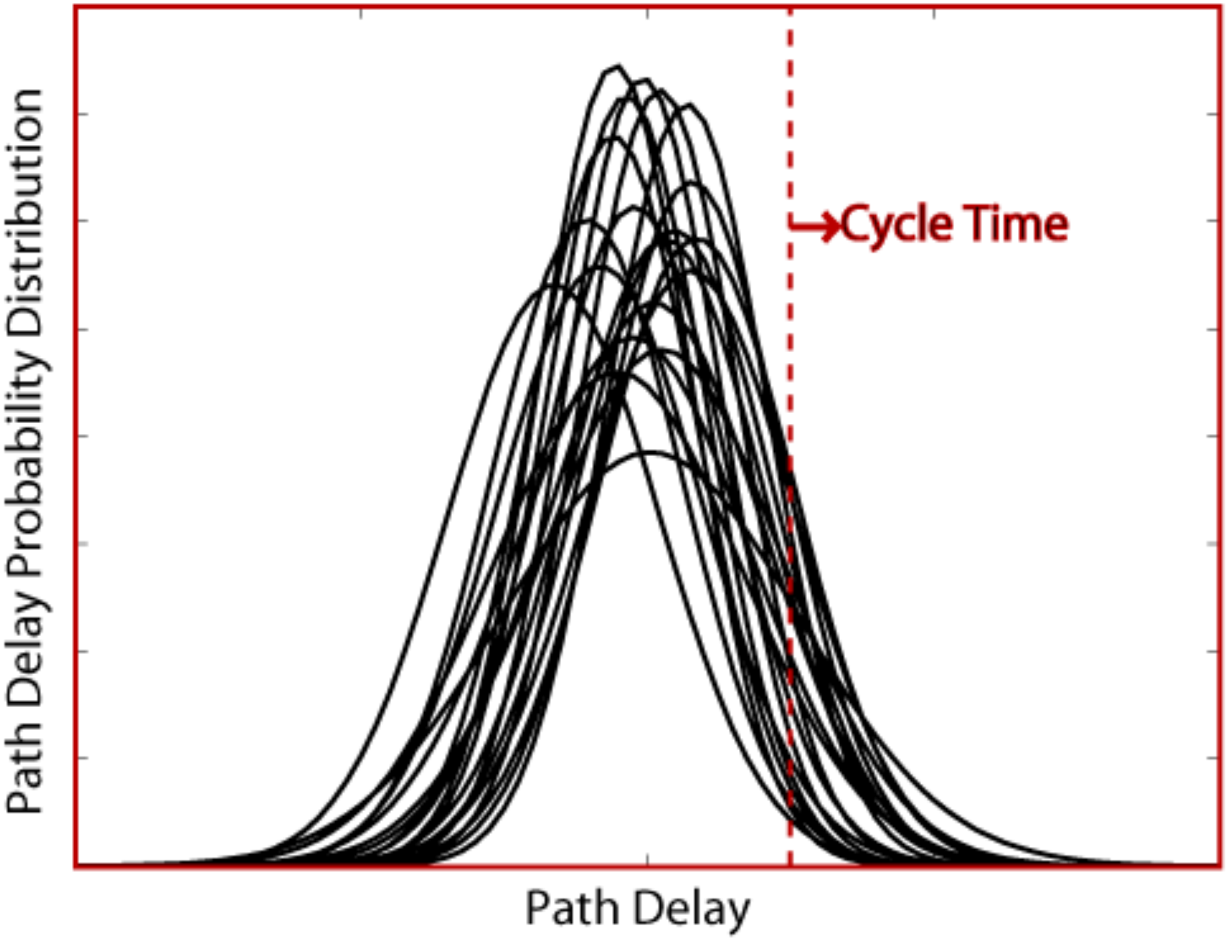

Adaptive Techniques for Dynamic Processor Optimization: Theory and Practice,
Wang and Naffziger (eds)., Springer, 2008

Circuit Delay Elements

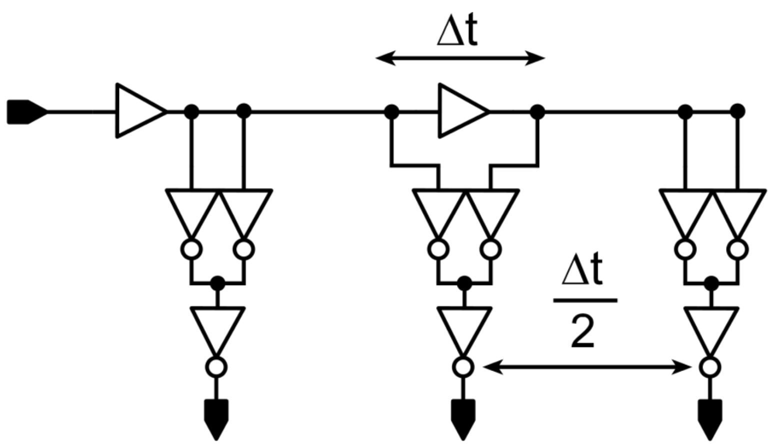

Interpolation

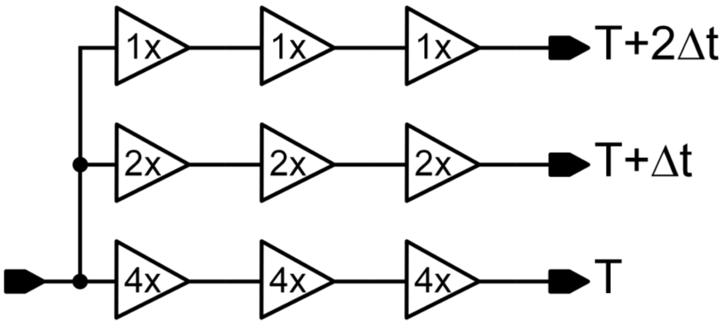

Variable Loading

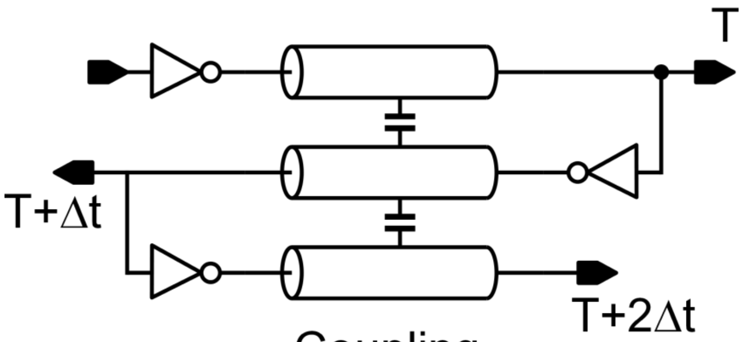

Coupling

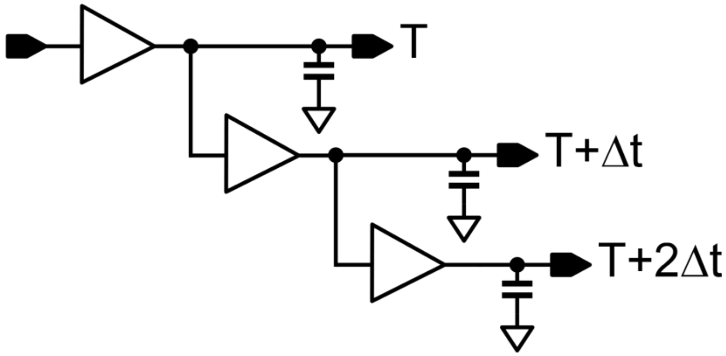

Capacitive Loading

Simplified Delay Line (Elmore)

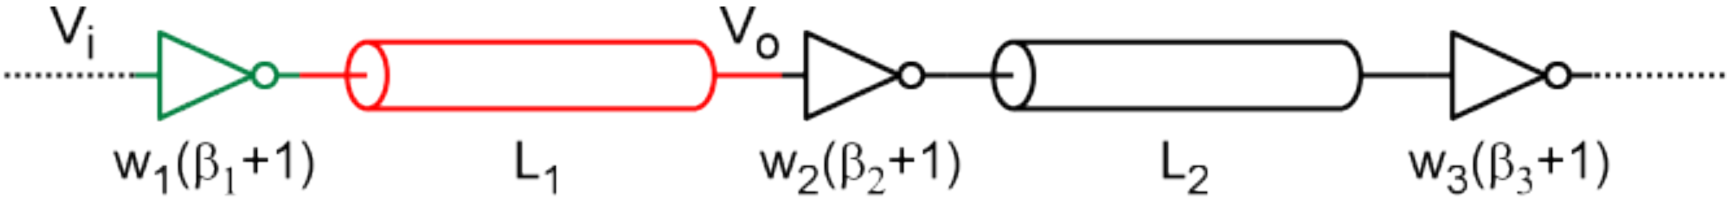

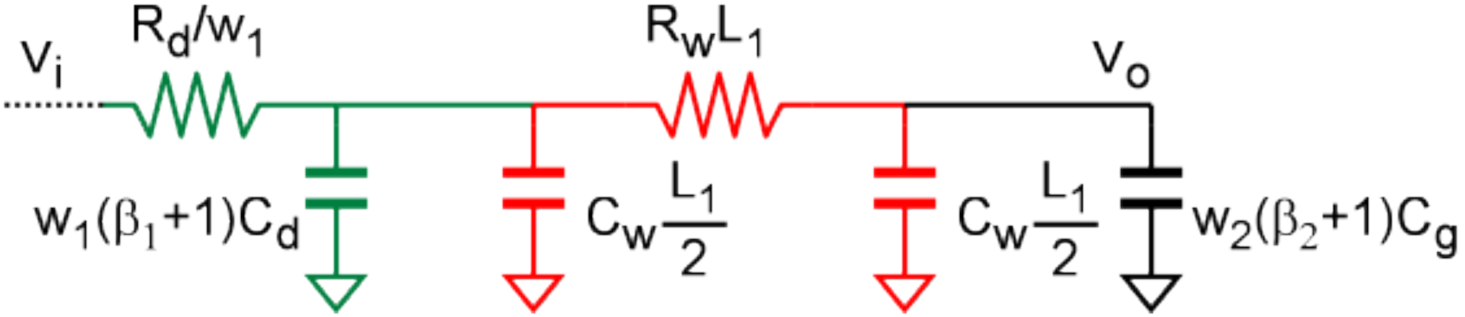

$$D_{path} = an\left\{\frac{R_d}{w}\left[w(\beta+1)(C_d)\right] + \left[\frac{R_d}{w} + \frac{lR_w}{\alpha}\right]lC_w + \left[\frac{R_d}{w} + lR_w\right]w(\beta+1)C_g\right\}$$

$$D_{path} = an\left\{\frac{R_d}{w}\left[w(\beta+1)(C_d + C_g) + lC_w\right] + l^2\frac{R_wC_w}{\alpha} + lR_w w(\beta+1)C_g\right\}$$

$$D_{path} = an\{D_{FET} + D_{RC}\}$$

Variation of Delay Sensitivity as Wire load Increases

$$D_{RC} = \gamma D_{FET}$$

$$D = an(1+\gamma)\frac{R_d}{w}C_{load}$$

$$S_x^y = \frac{x}{y}\frac{\partial y}{\partial x}$$

$$S_\gamma^D = \frac{\gamma}{1+\gamma}$$

$$S_{R_d}^D = 1$$

- If ratio of wire to gate delay is fixed, sensitivity is 1 to environmental changes. Can use any mix of wire/gate delay.
- Ratio will not be fixed (wire delay becomes different ratio as environment changes)
  - Highly voltage dependent
  - Wire dominates at high voltage

Wang, 2008

Delay Sensitivity to FET/Wire Ratio γ

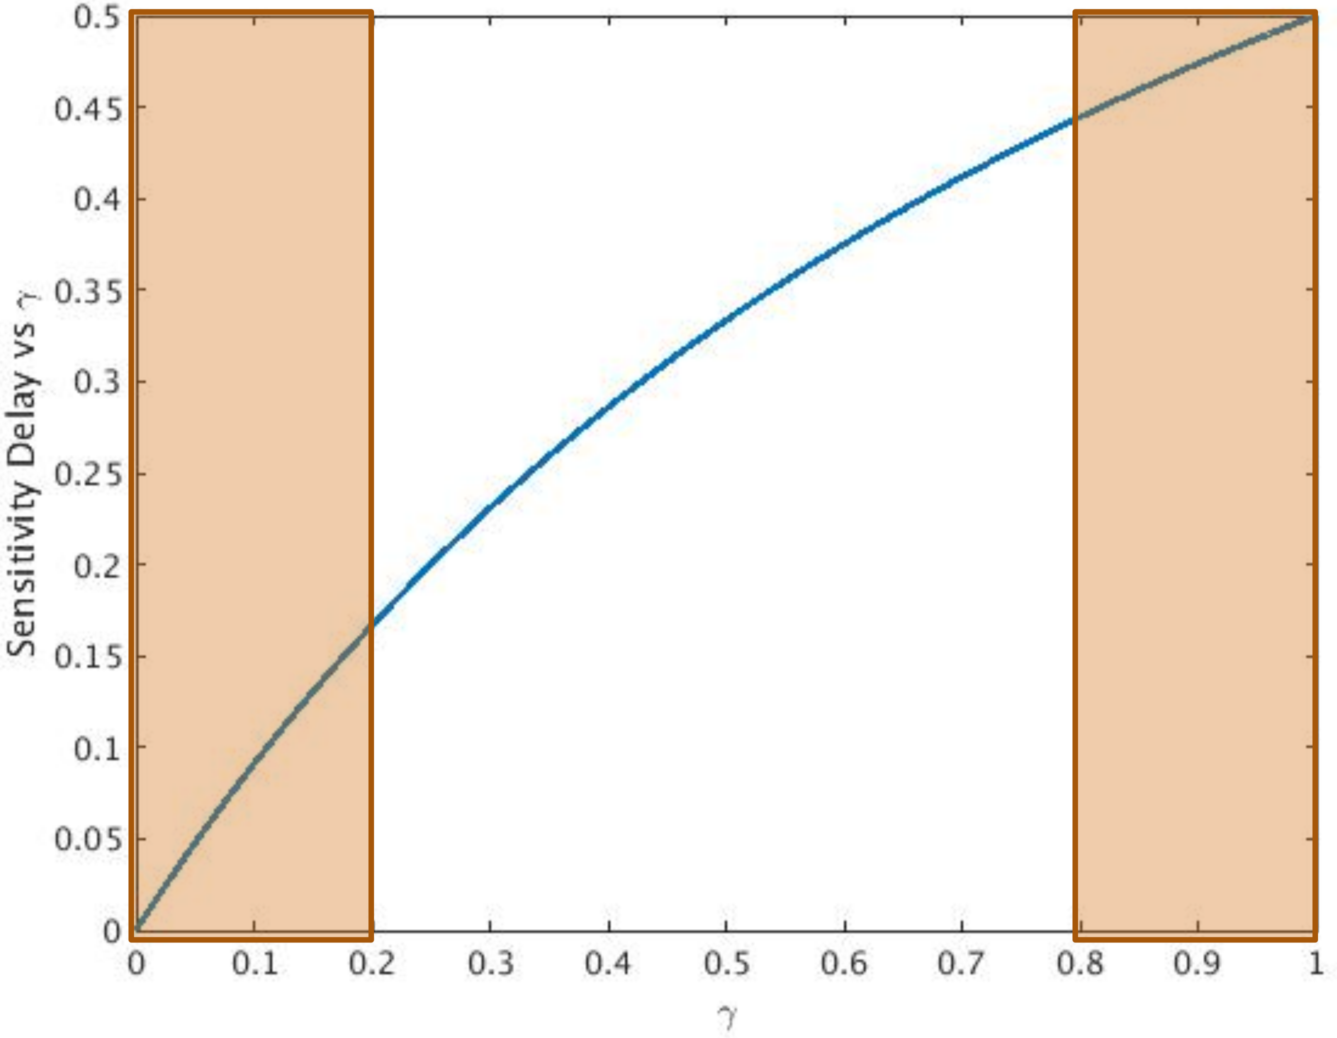

Prohibited by design rules.

Critical Path Synthesis

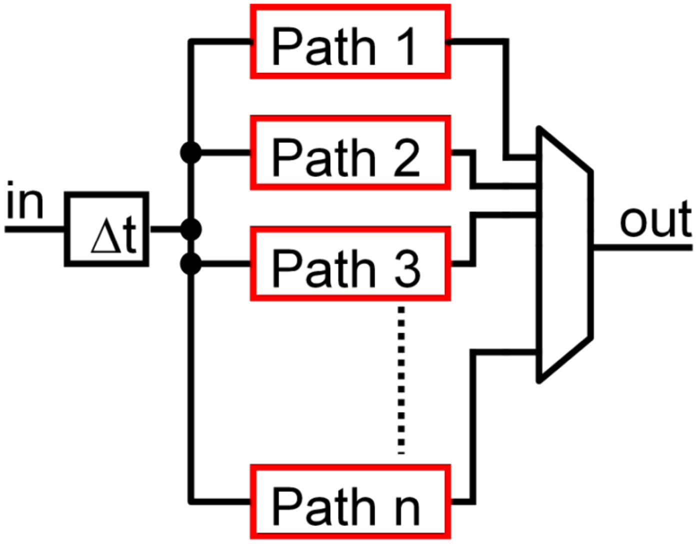

Parallel paths

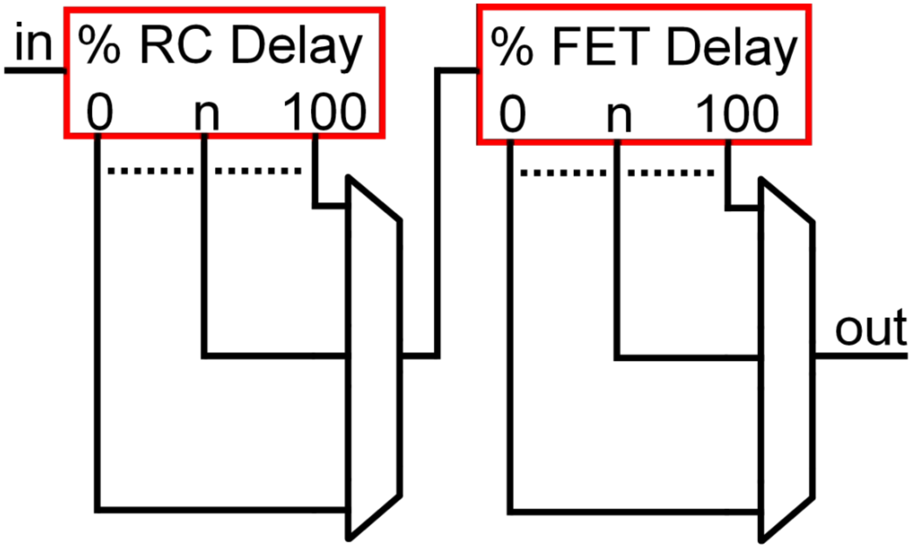

Serial paths

Wang, 2008

CPM Path Synthesis

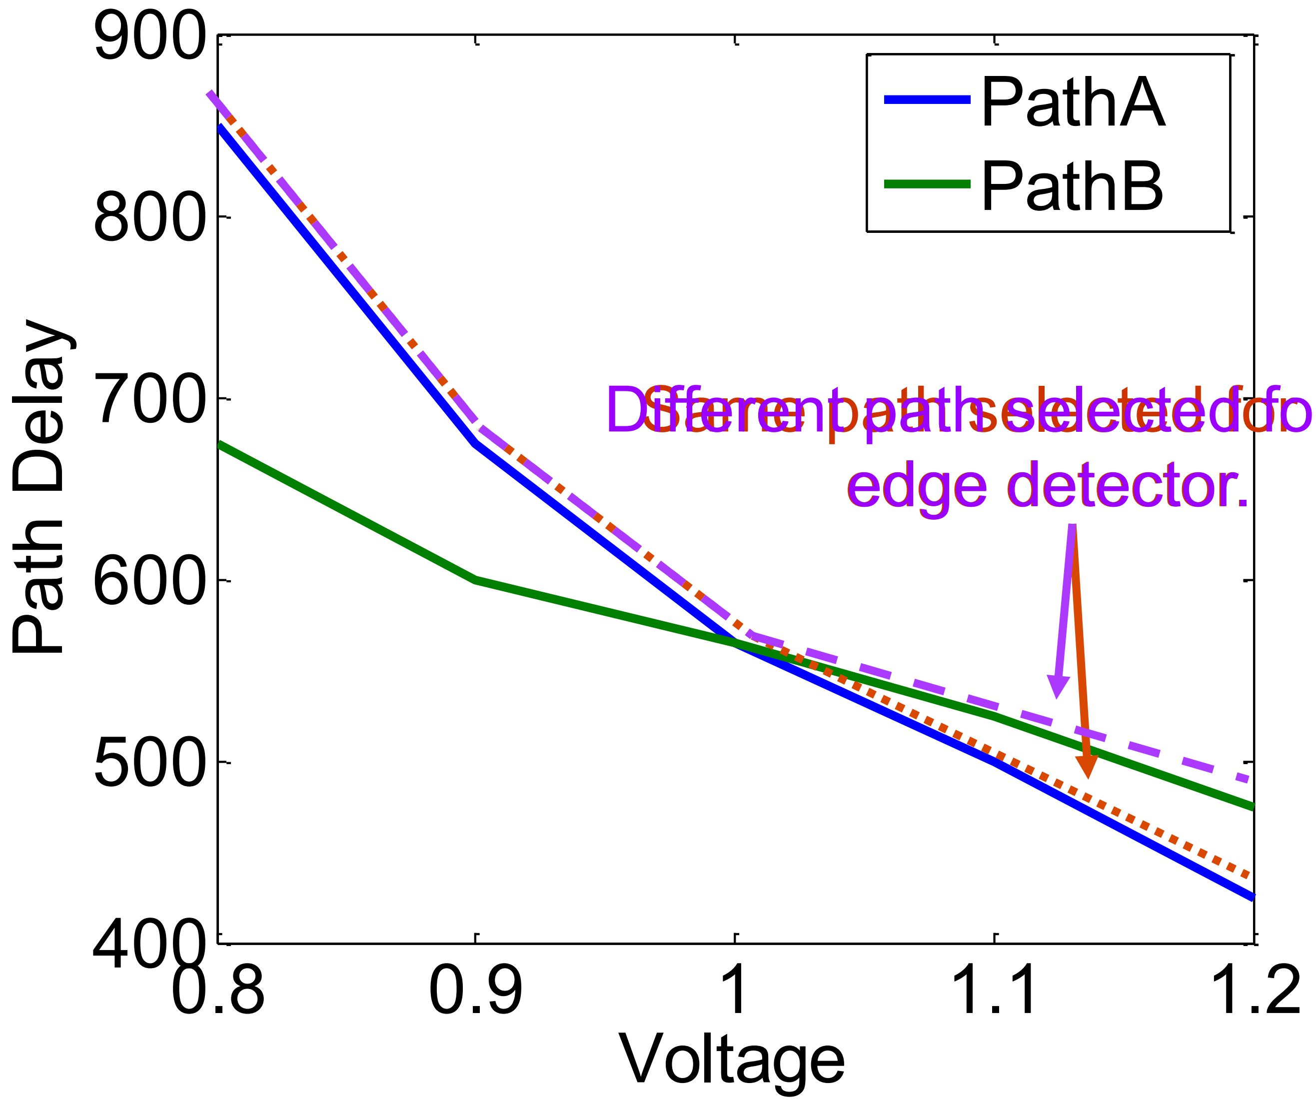

Different paths selected for each edge detector.

Same path selected for each edge detector.

What is the Appropriate Path Configuration?

- Fixed variation eliminated during calibration
- How closely does the monitor have to track critical path delay?
- How sensitive should the monitor be?
- Random and noise generated variation causes timing failures
- Dependent on the number of gates in a critical path
- FET delay behaves consistently over device type and topology
- Design rules limit degrees of freedom in potential critical paths
- Count the number of gates

Time-to-Digital Converter Design Considerations

- Metastability
  - Reduce edge rates
  - Add a secondary latching stage
- Rising/Falling capture time differences
- Pulsed versus non-pulsed clocking
- Capture window width
- Capture resolution
- Robustness

Simple D Flip-Flop Phase Detector

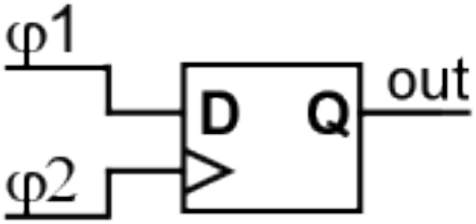

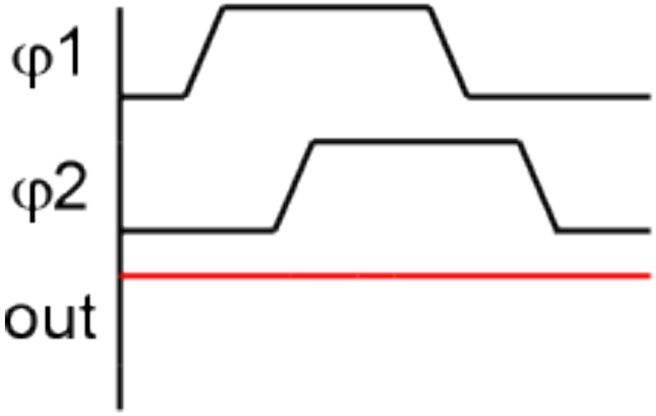

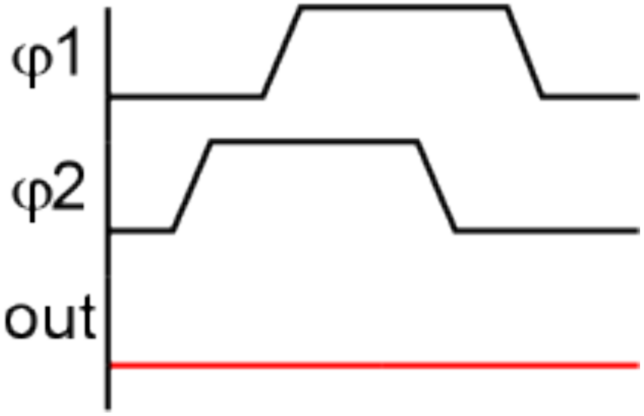

Wang, 2008

Phase-Frequency Detector

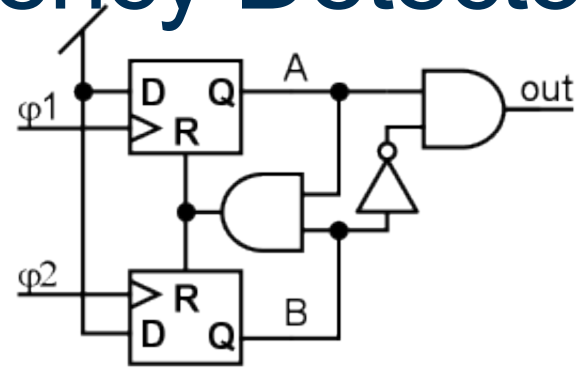

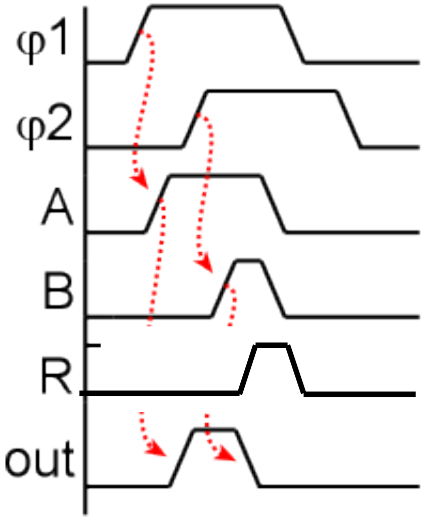

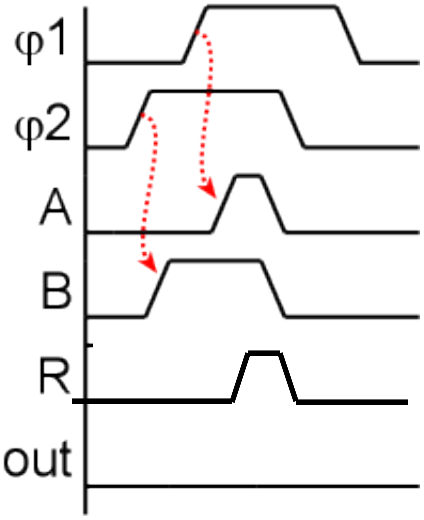

Wang, 2008

Delayed Phase Converter

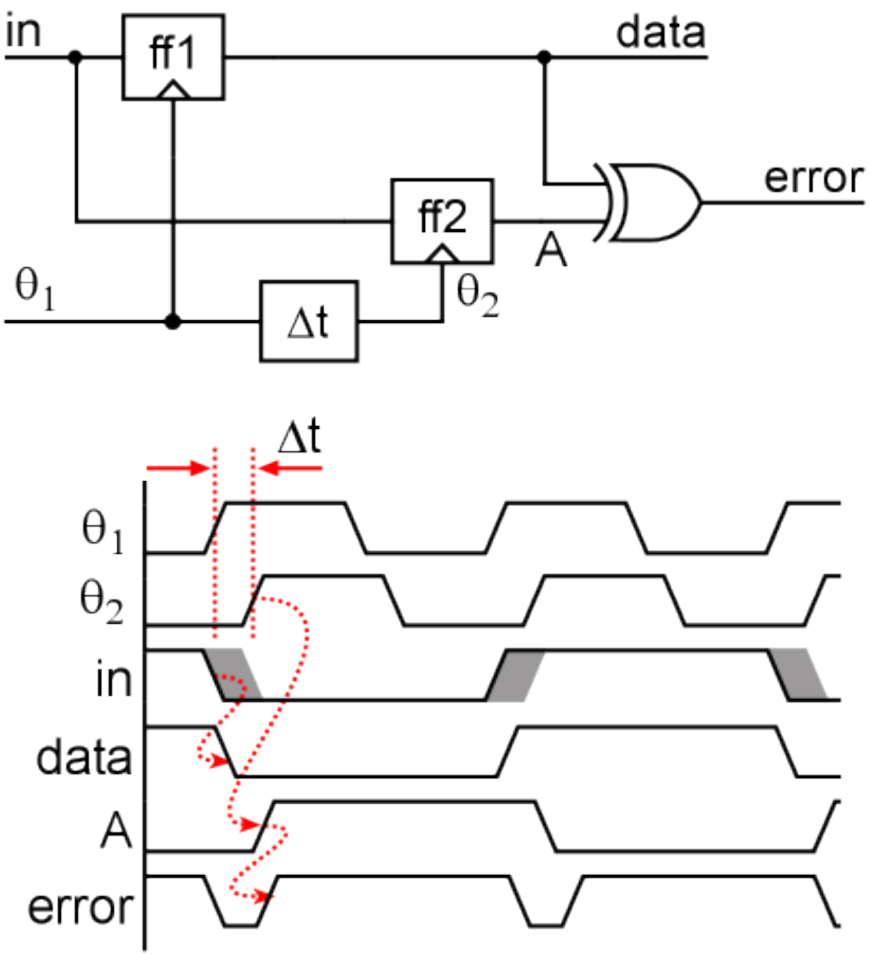

Austin, et. al.,

Computer, Mar. 2004, pg. 57

Multi-Bit Time-to-Digital Converter

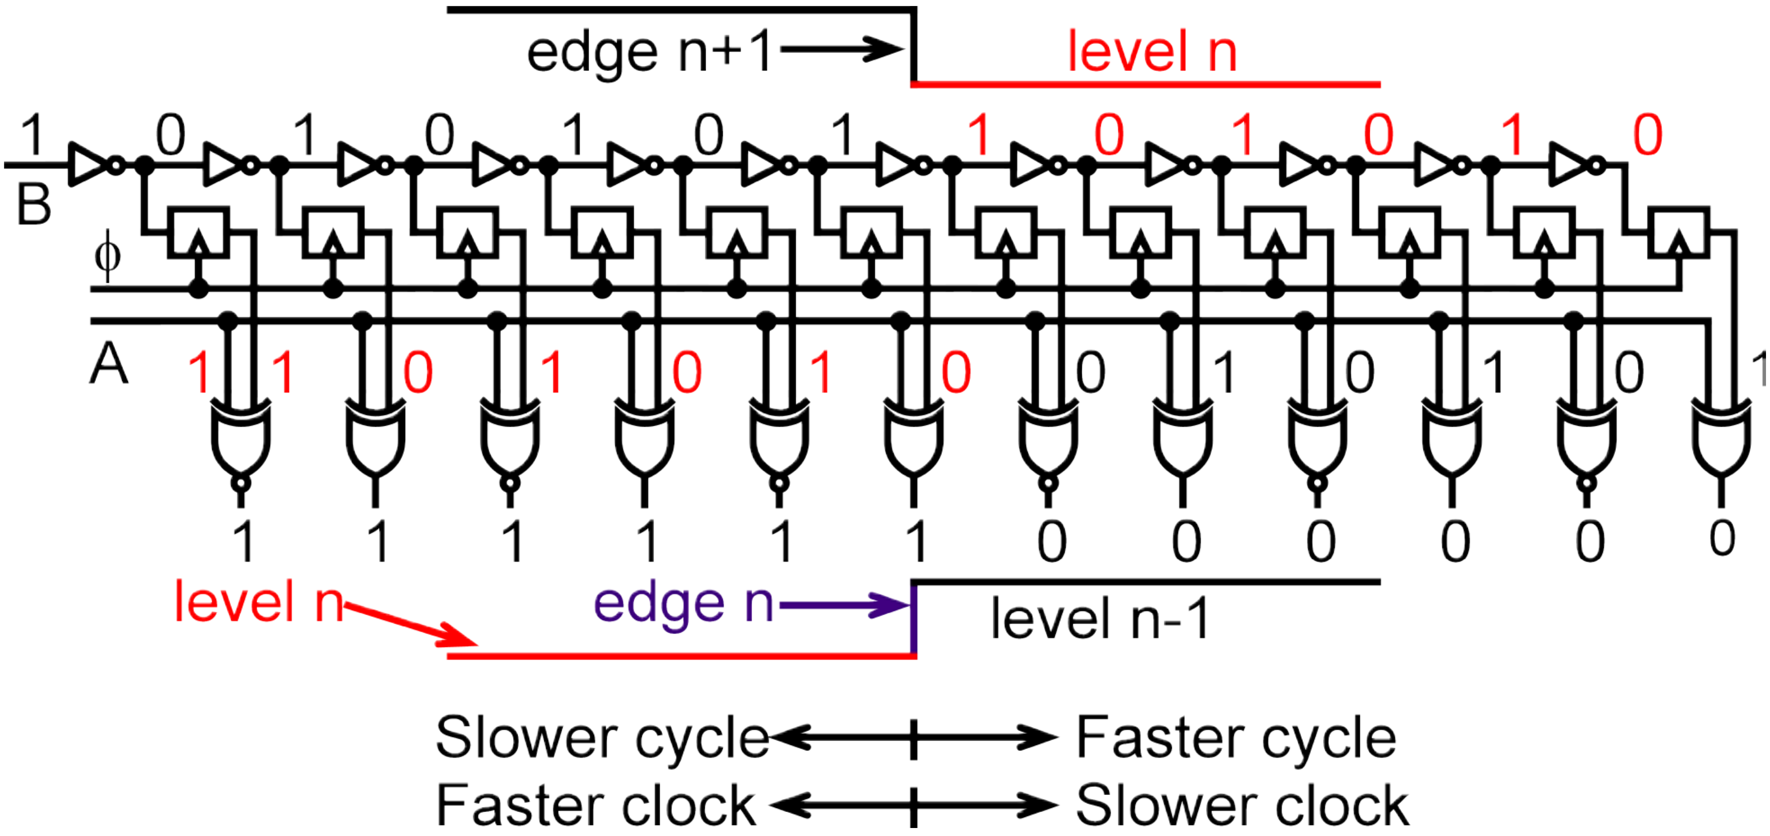

Drake,
ISSCC, 2007

Multi-Bit Output Variations

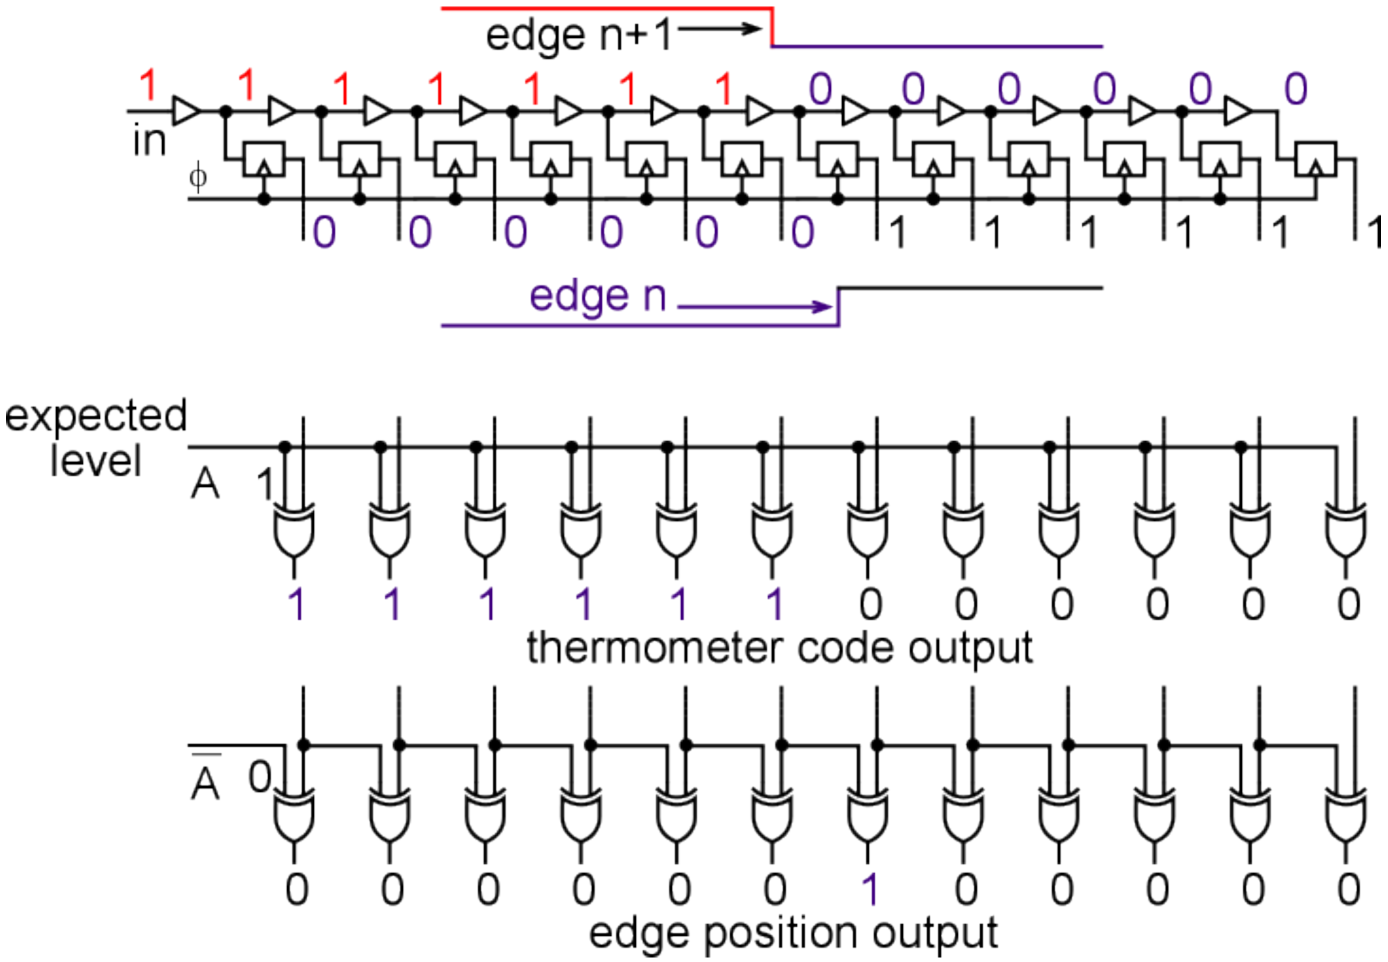

Wang, 2008

Interpolating Delay Line

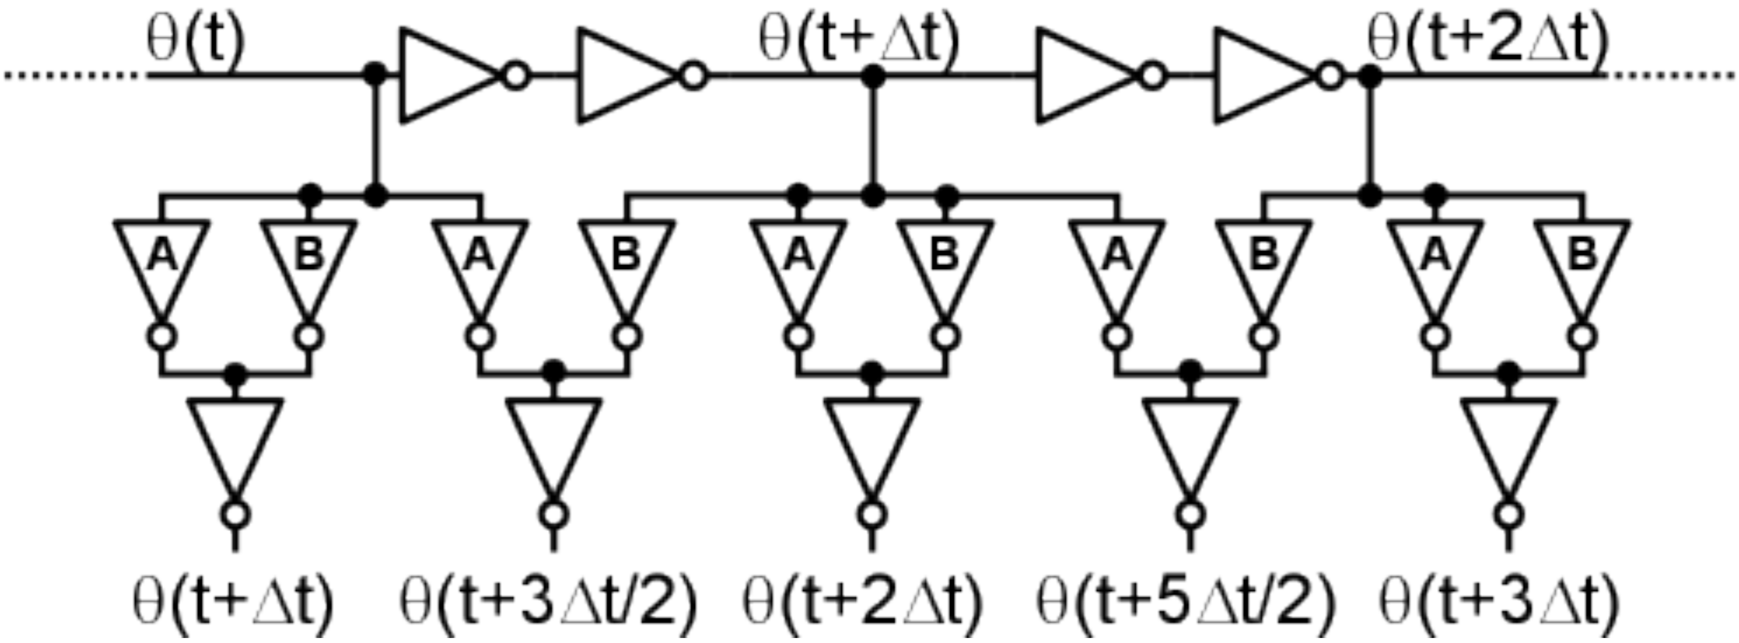

$$0.6 = \frac{W_A}{(W_A + W_B)}$$

Garlepp et. al.,

JSSCC, May 1999

POWER6 CPM Block Diagram

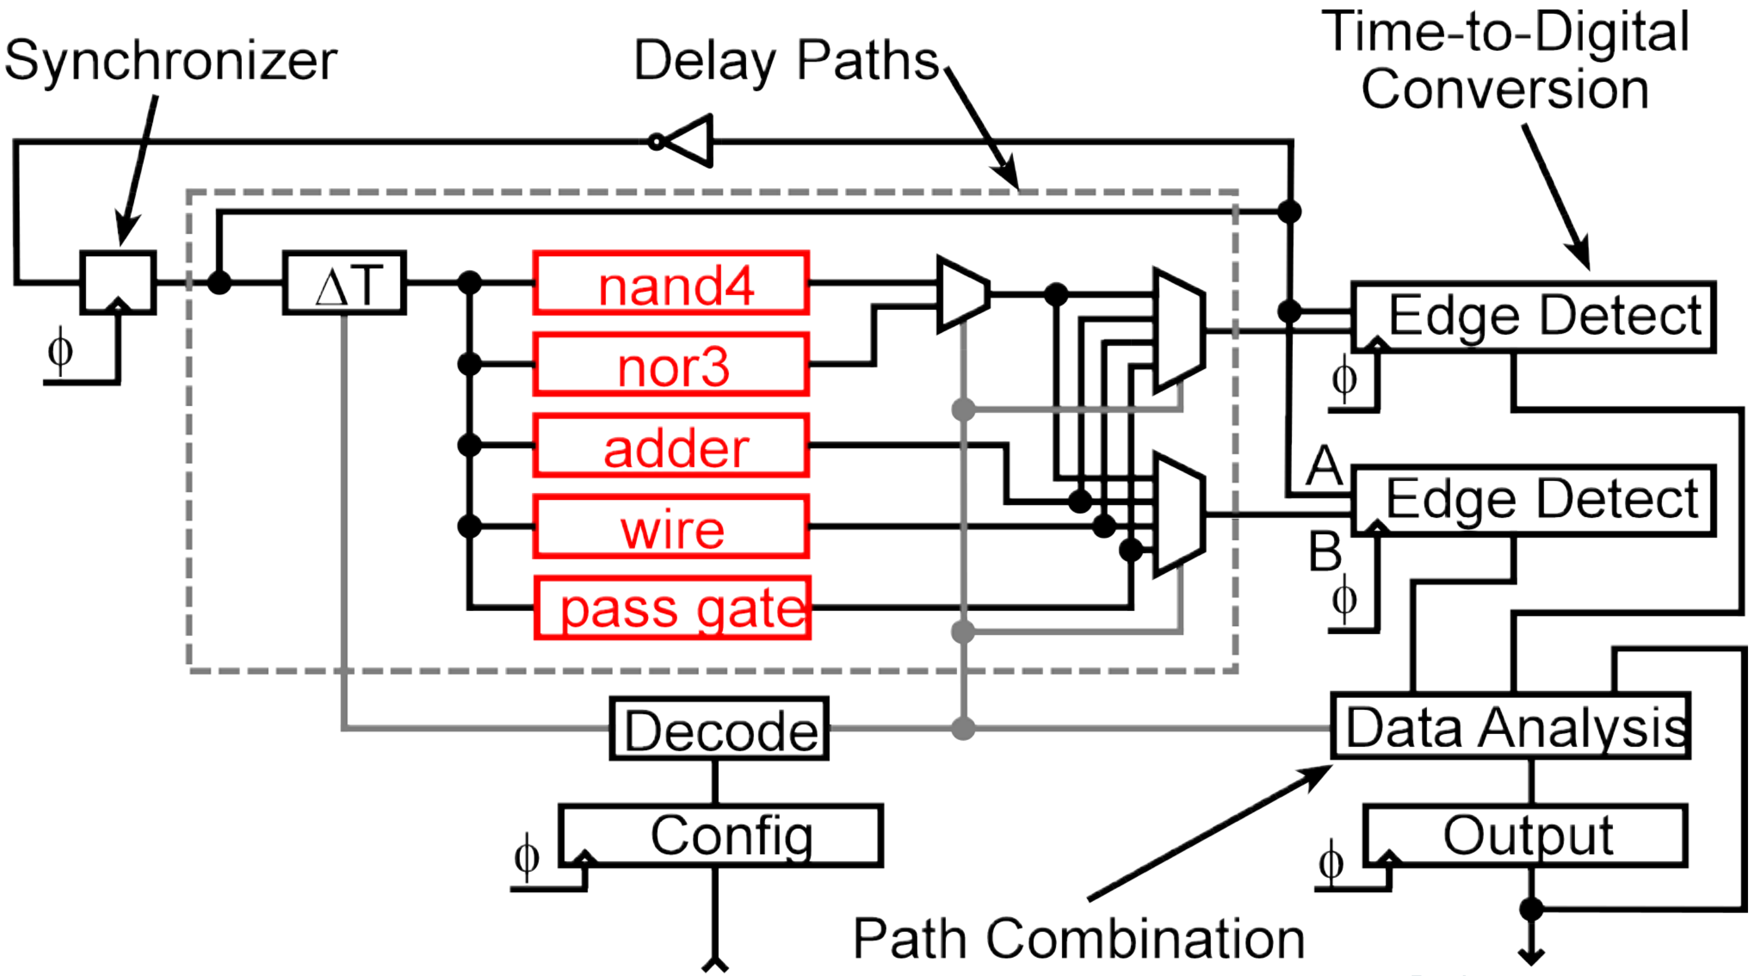

Drake,
ISSCC, 2007

POWER6 CPM Characteristics

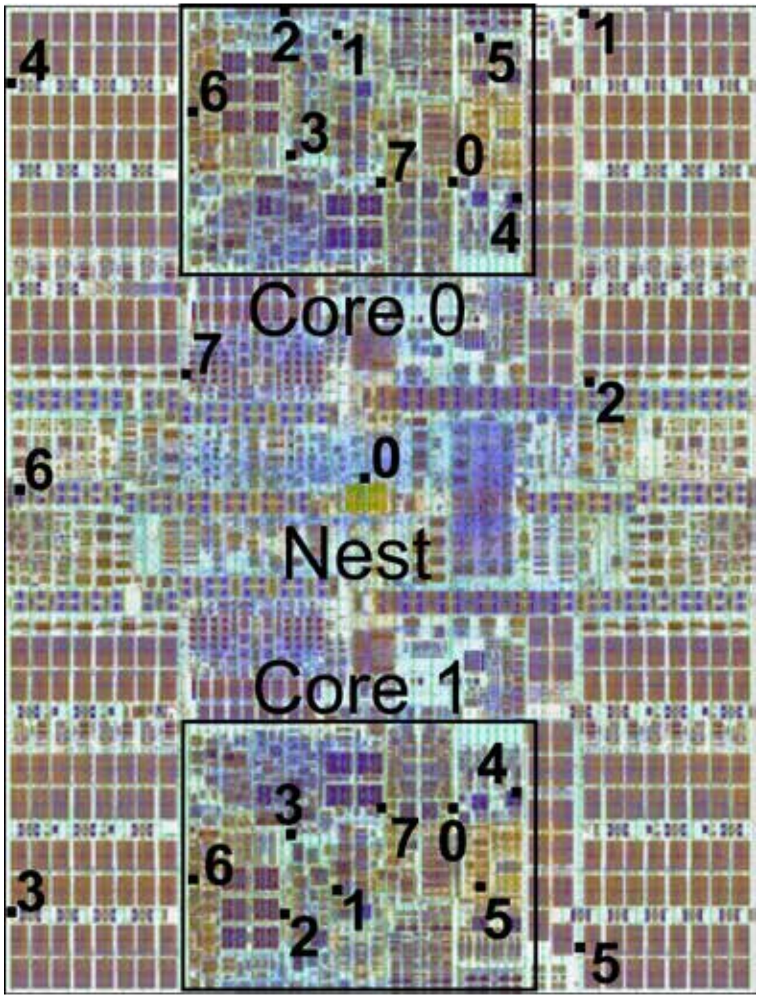

- 24 CPMs on chip
- 2 point calibration
- 12 bit output
- 20mV/bit AC noise
- 10mV/bit DC noise
- 10°C/bit temperature
- 1 sample/cycle
- Core CPM is $90 \times 36 \mu m^2$
- Nest CPM is $90 \times 48 \mu m^2$

Drake, *ISSCC*, 2007

Core0 CPM Behavior Under Cycle-Repeatable Workload

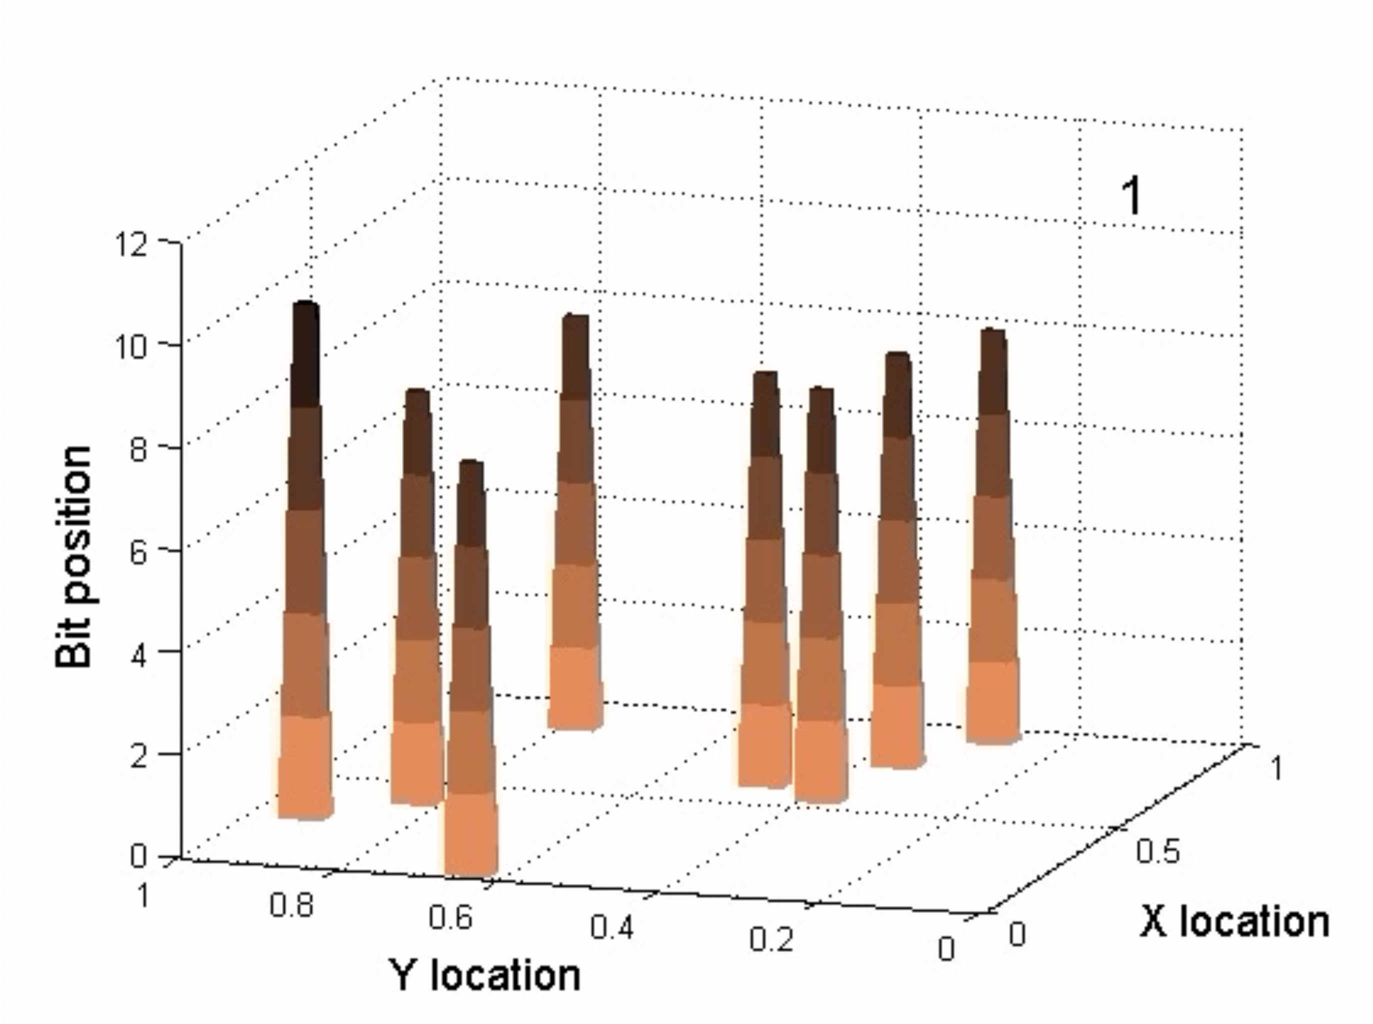

POWER7+ Critical Path Monitor Overview

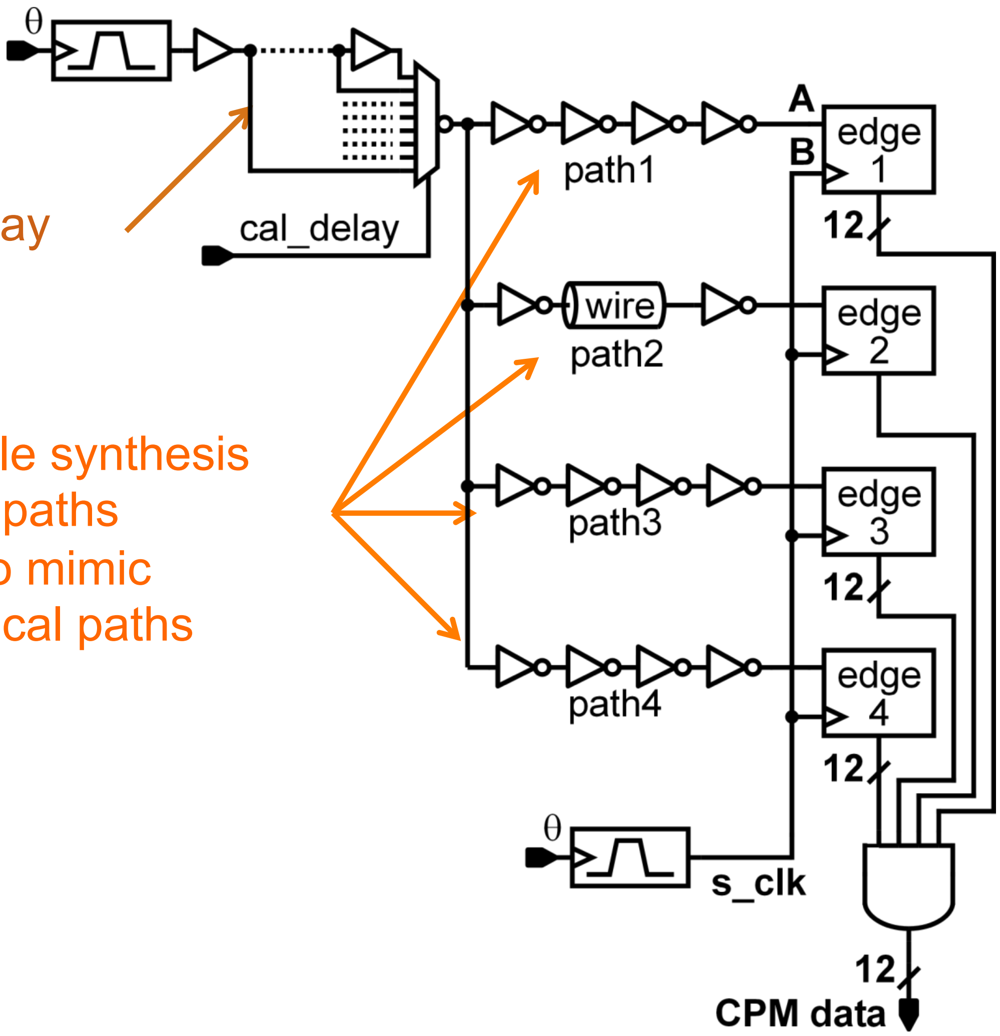

Calibration Delay

Multiple synthesis paths to mimic critical paths

Power7+ Controller Diagram

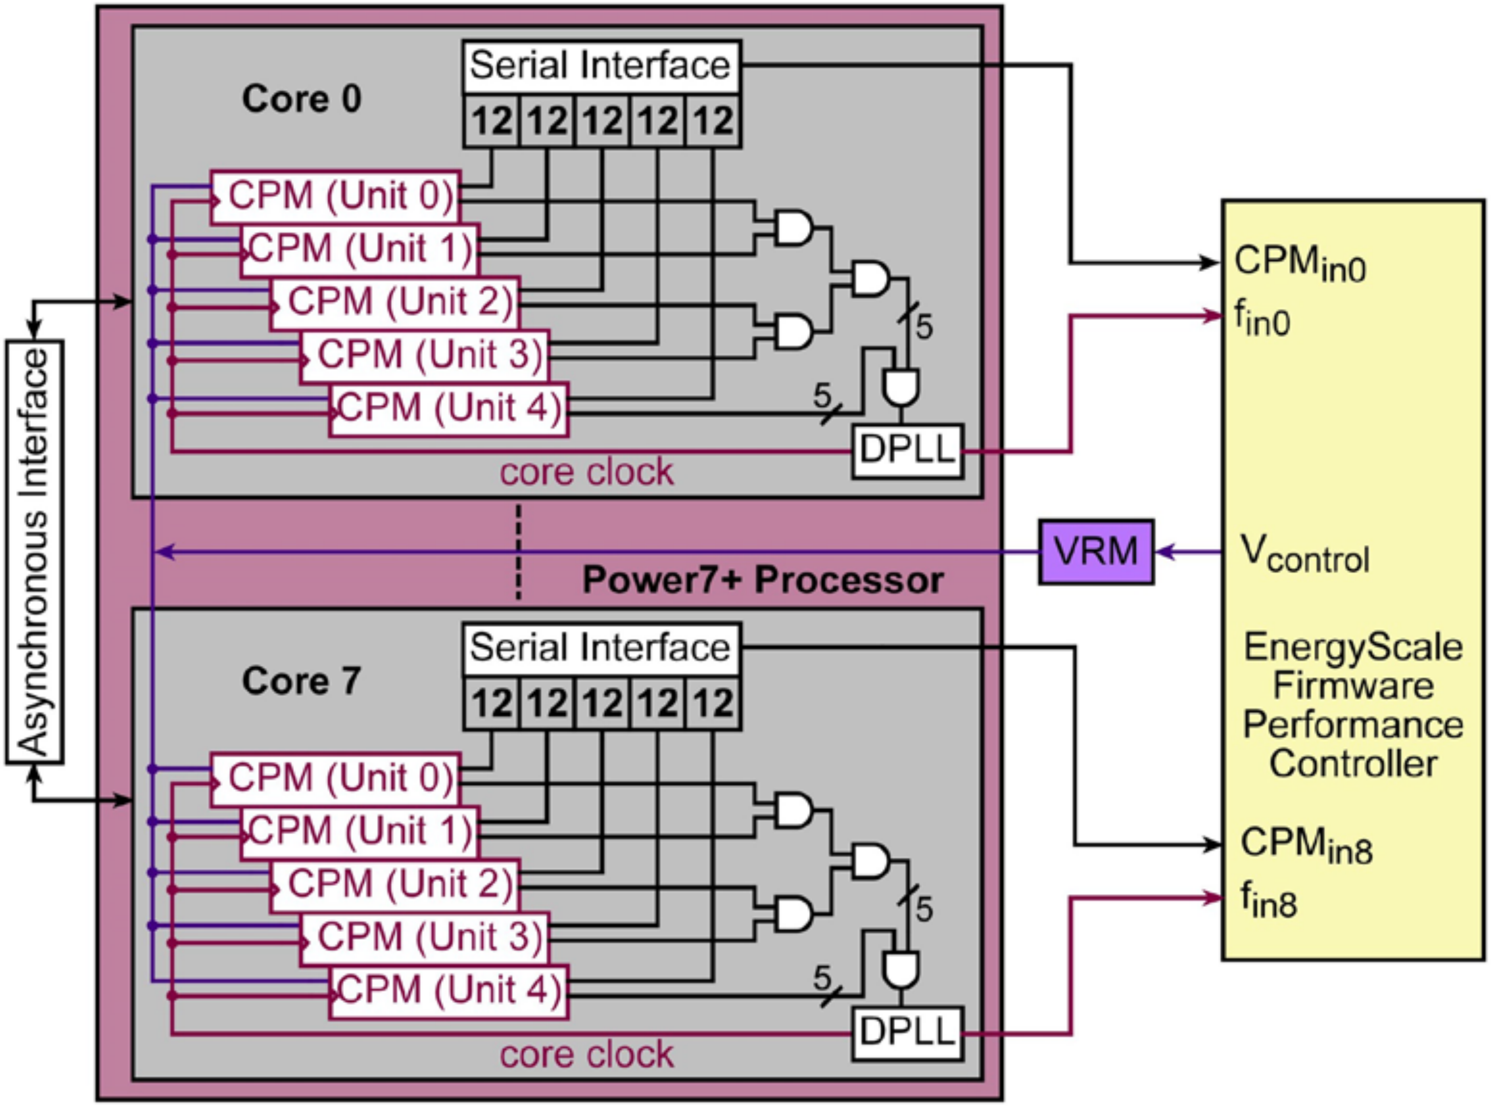

Critical Path Monitor Block Diagram

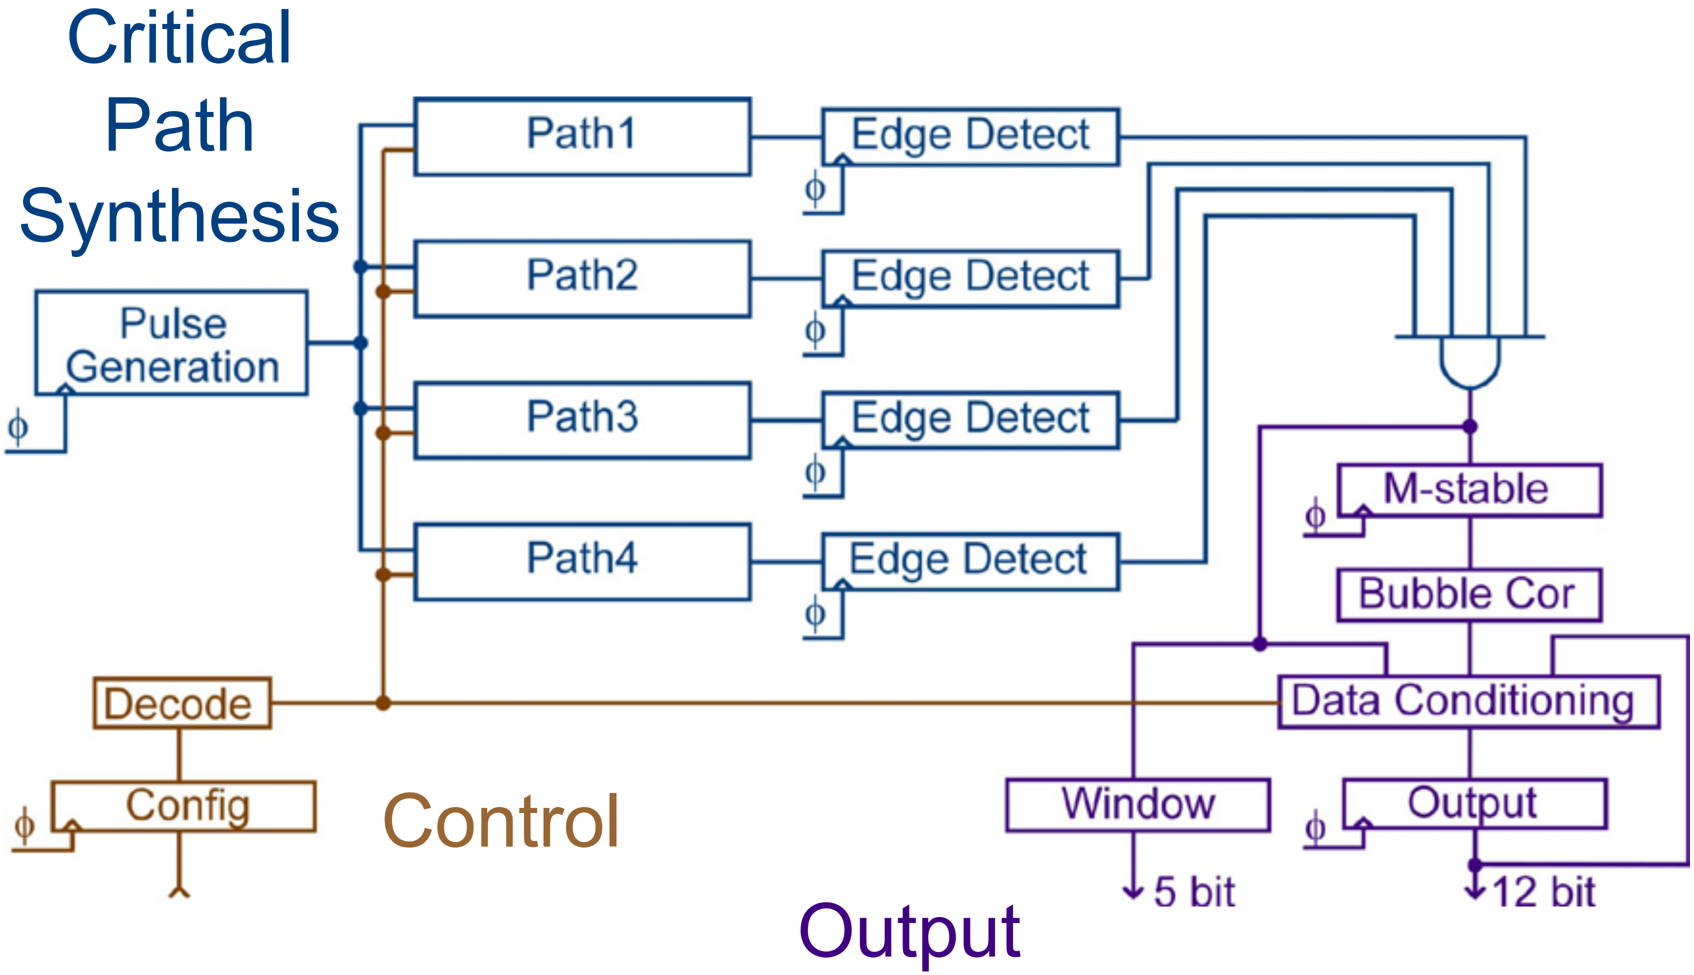

Path Synthesis Block Diagram

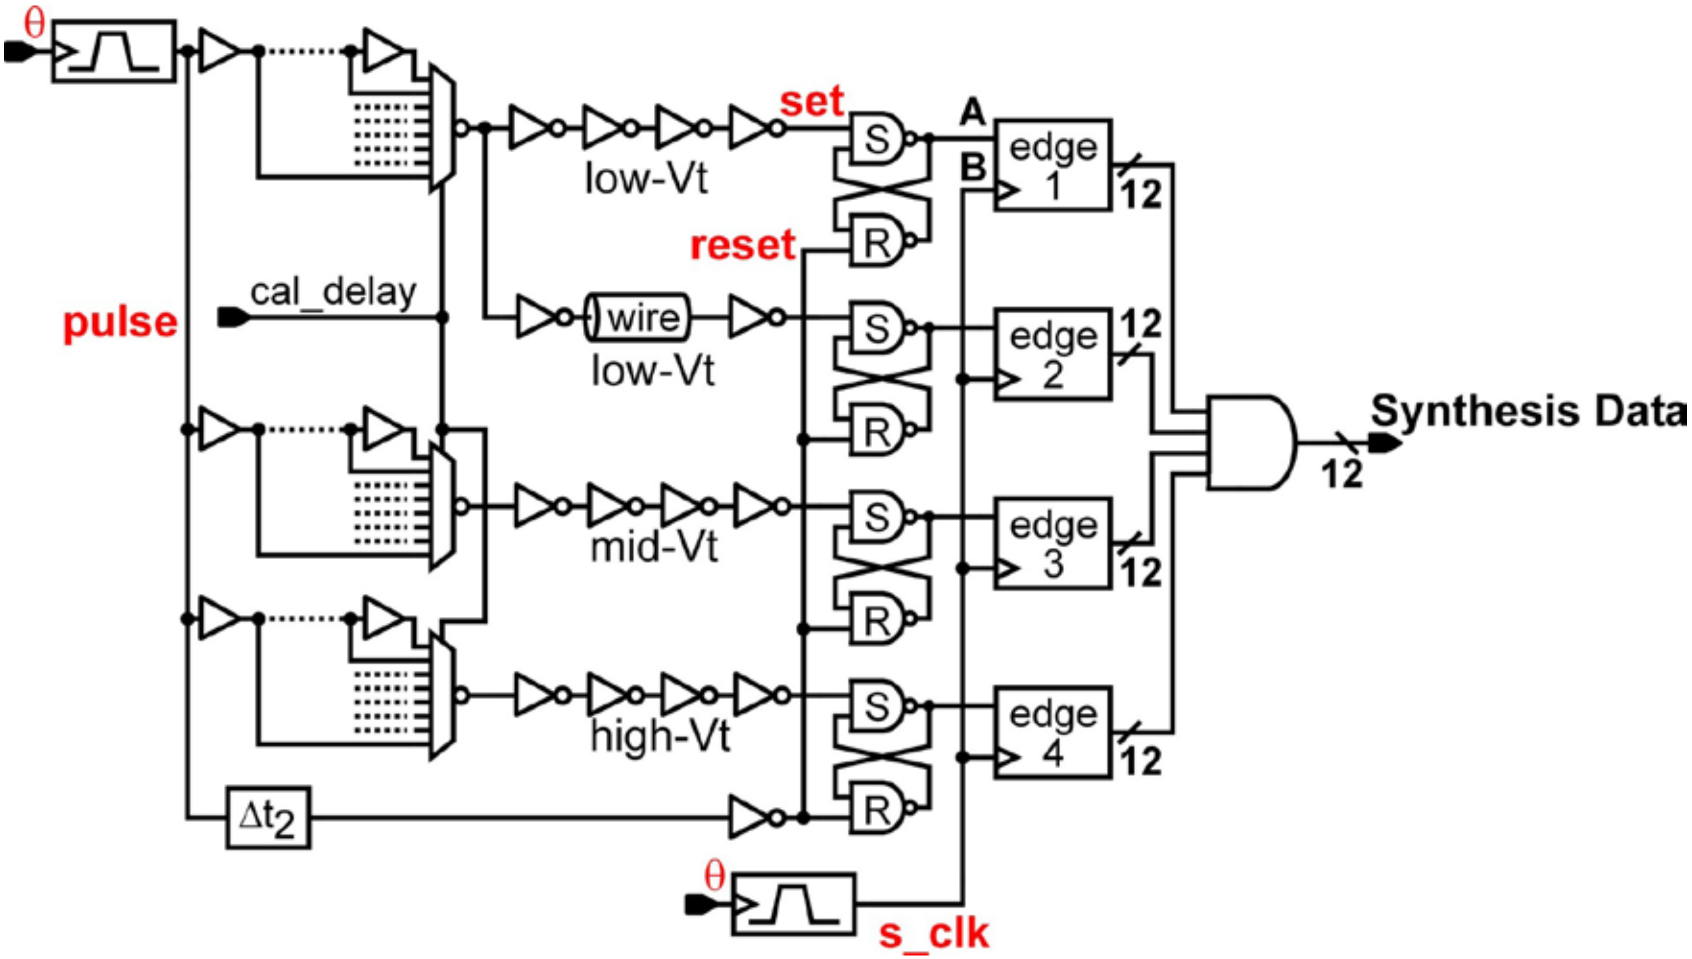

Synthesis Path Pulse Shaping

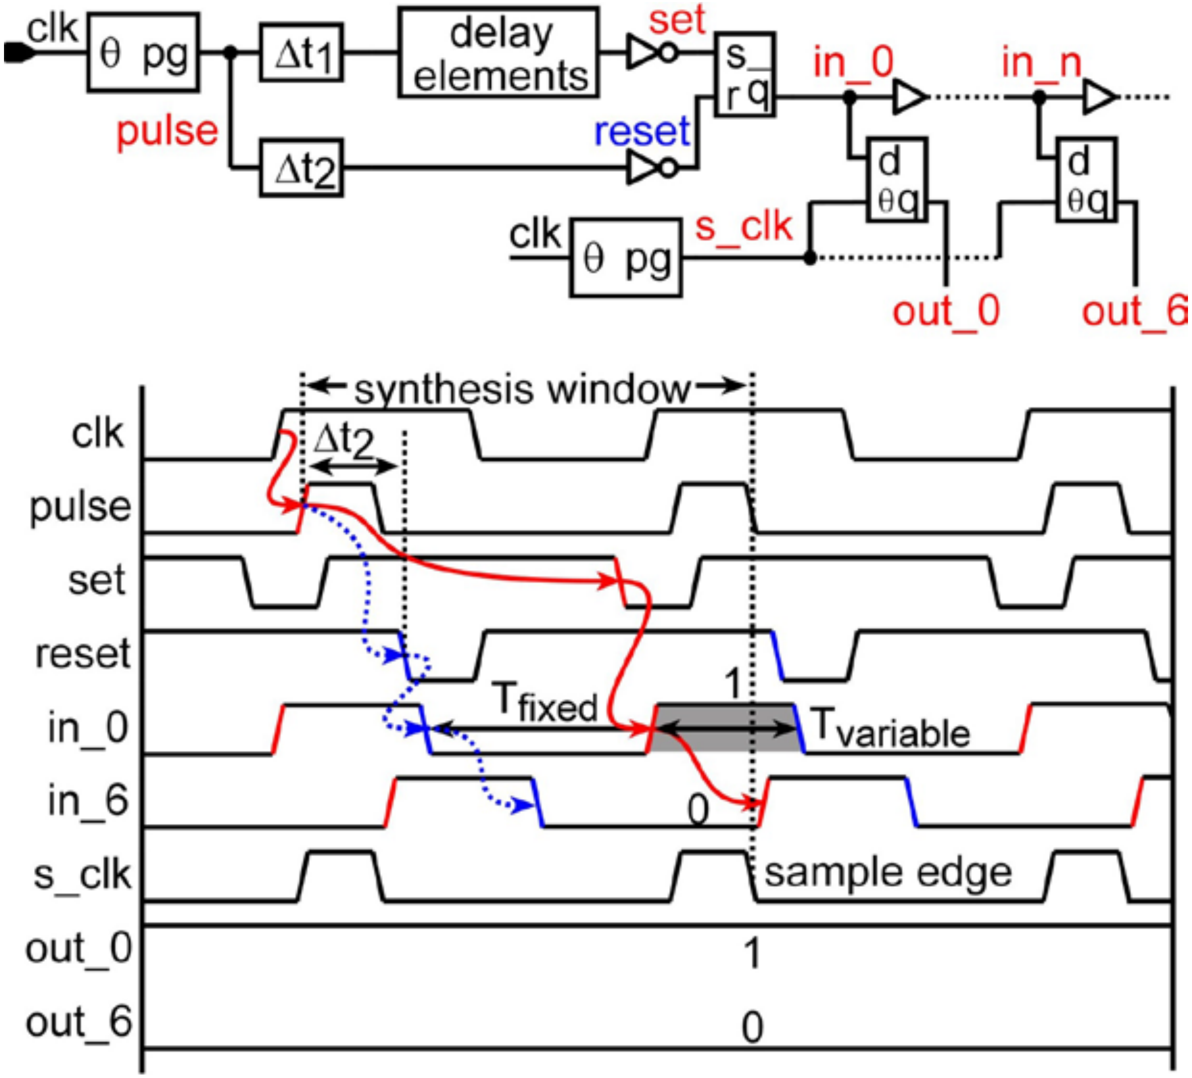

High-Frequency Limit

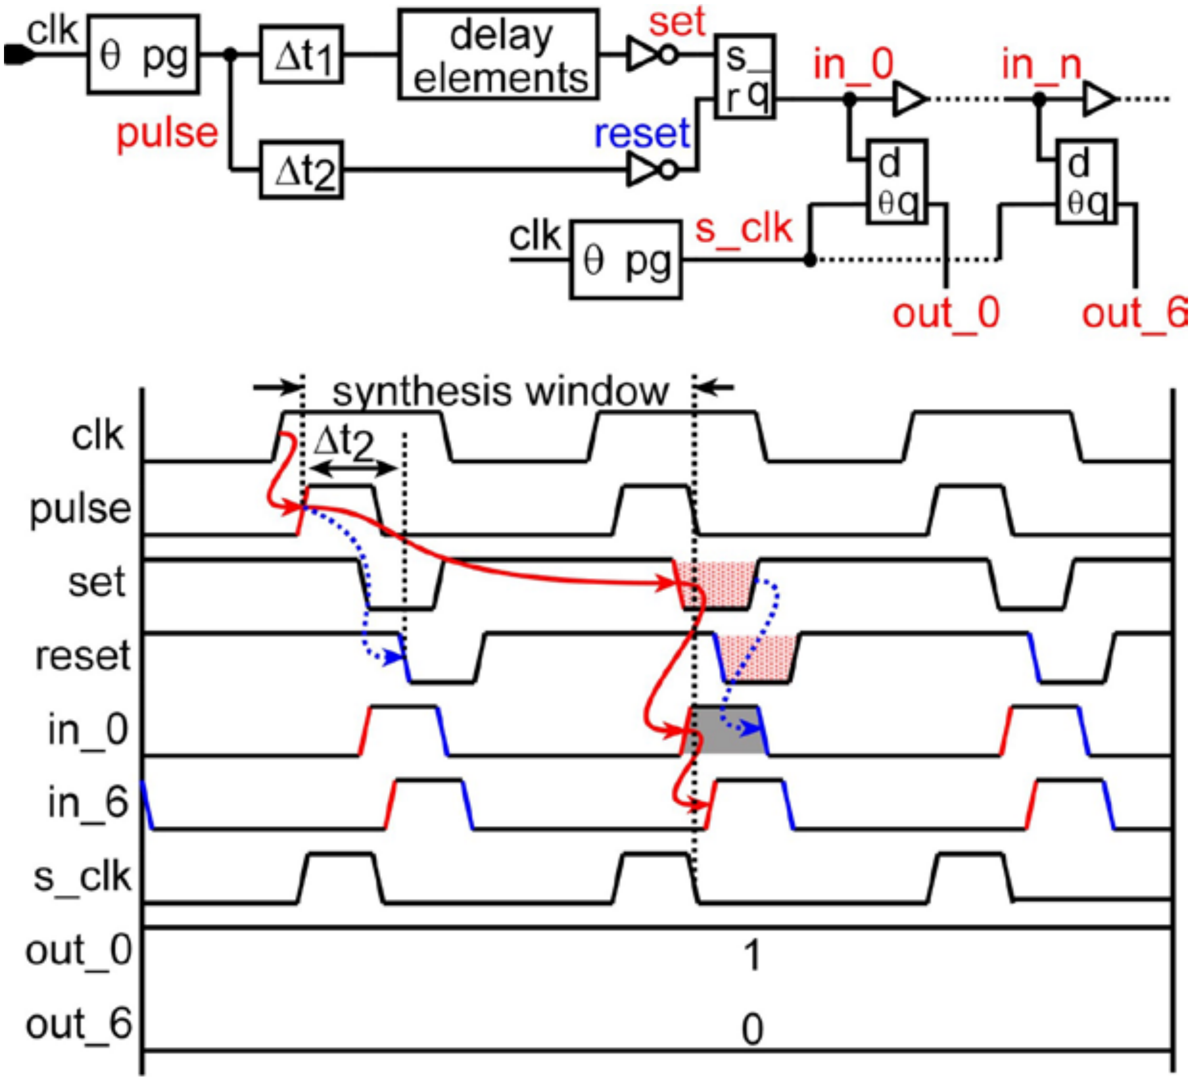

- Pulses Overlap at High Frequency
- Reset pulse loses control
- Designed to be safe up to 10% above maximum frequency over process and operating range
- No lower frequency limit

Edge Detector/Fine-Delay Block Diagram

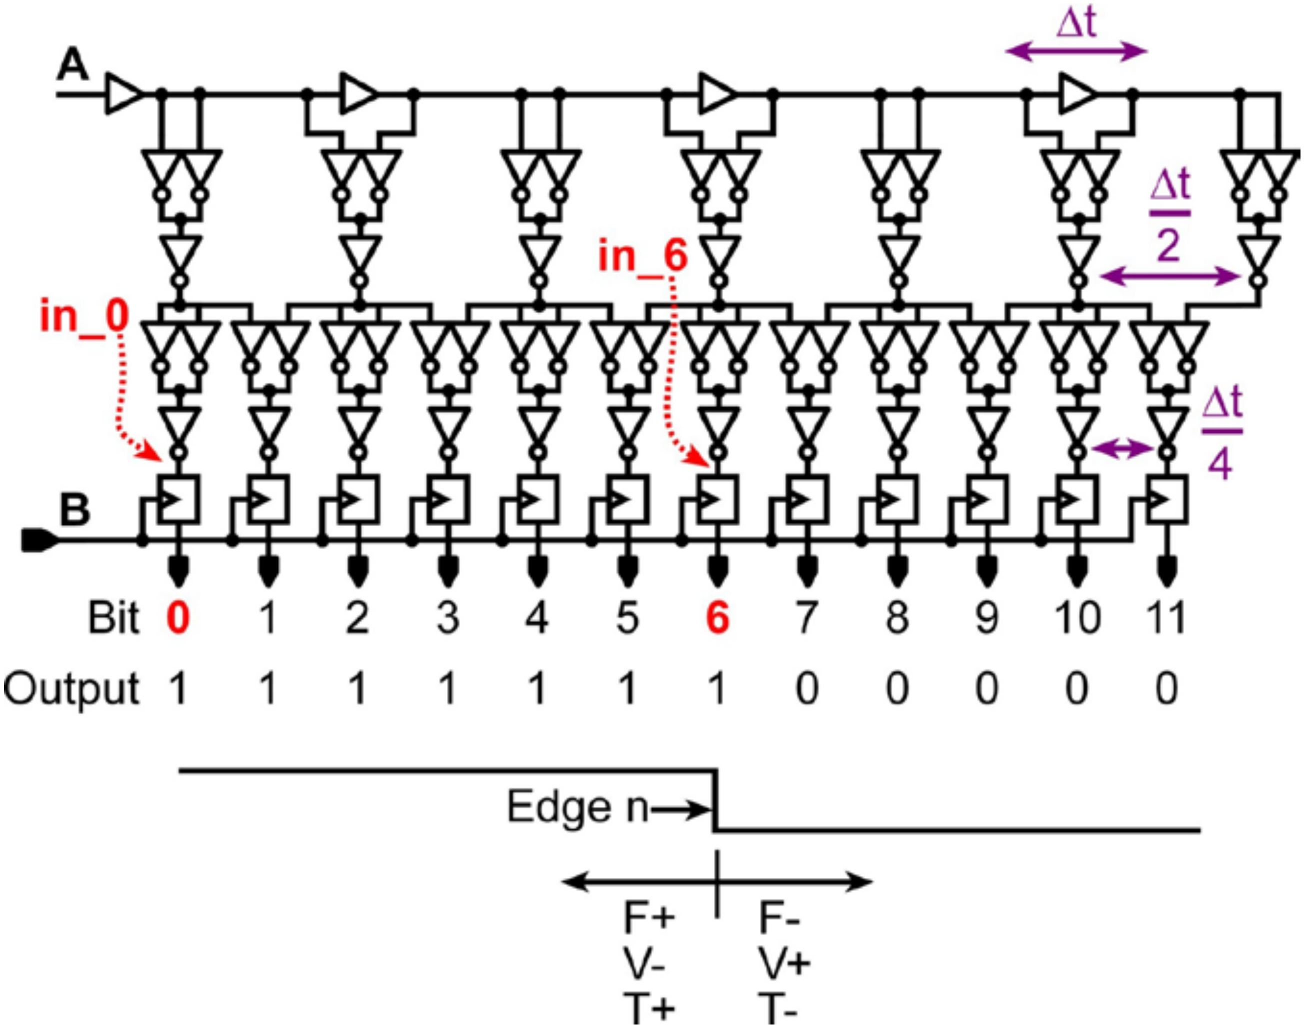

Calibration Delay

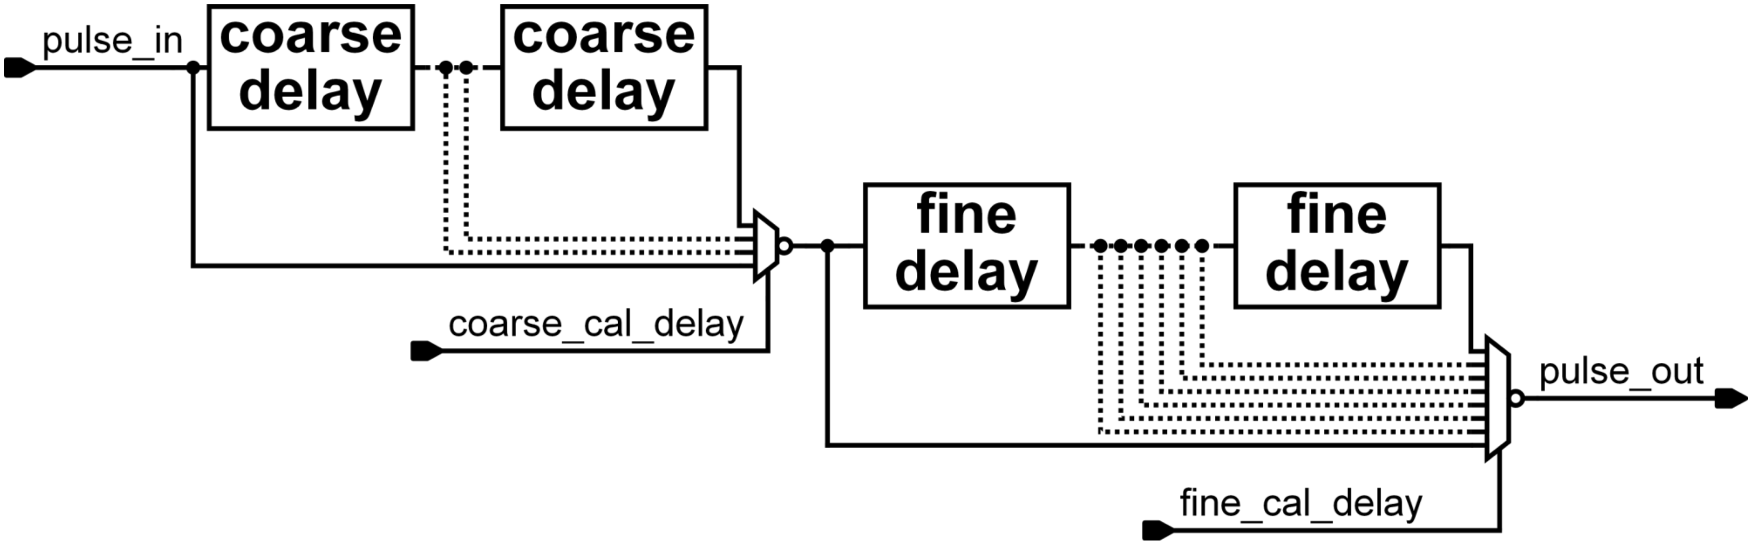

- Coarse delay block equal to sum of fine delay blocks
- Muxes combine delay into continuous range
- Uses delay elements shown earlier

Tuning Methodology

- Hierarchical Simulation
  - Fine delay lines only
  - Complete path without control
- Multiple runs
  - 4 paths
  - 2 guardband frequencies
  - 3 operating points: powersave, nominal, and turbo
  - multiple process corners

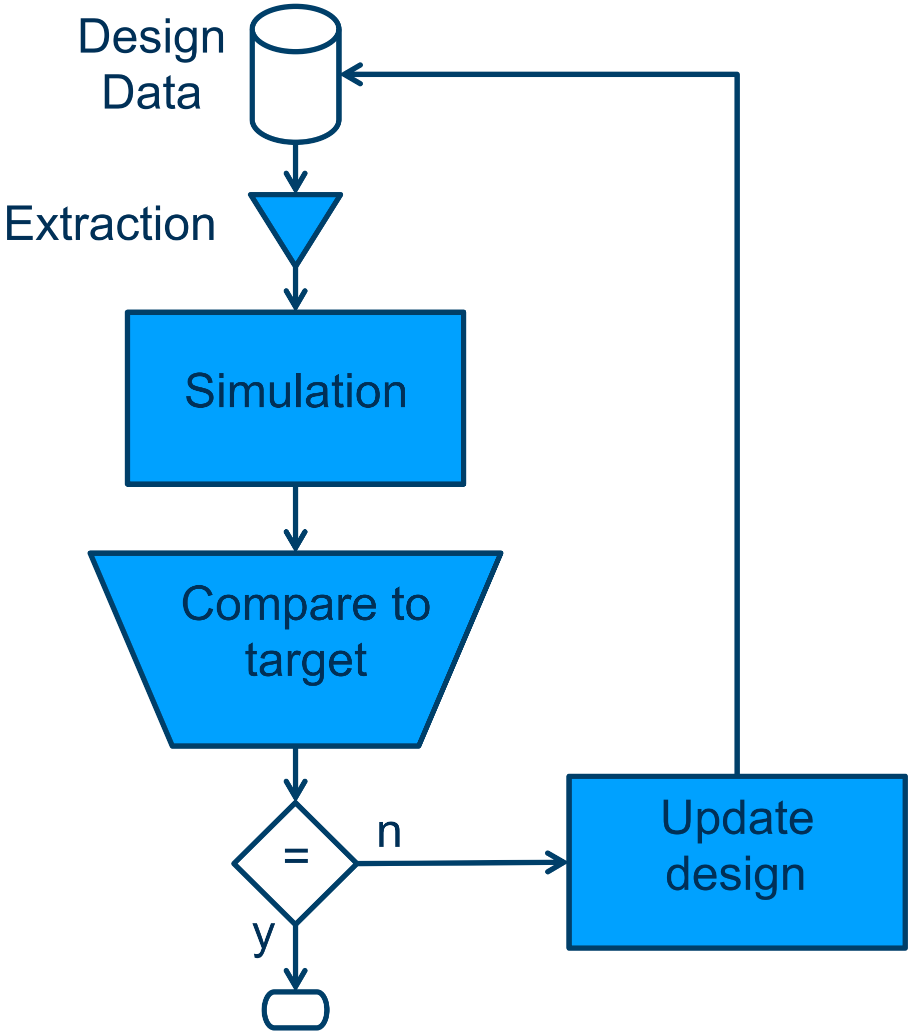

Extraction and Simulation

- Hierarchical
  - Control: 10724 transistors
  - Custom synthesis block: 5257 transistors
- Extraction
  - Full coupling for power and noise analysis
  - Reduced coupling for timing
- Simulation
  - Spice
    - Slow (6hrs, 34mins) but flexible
    - Accuracy dependent on extraction
  - Static timing
    - Fast (50 minutes)
    - Very accurate at 2 points: early and late mode timing

HW1: Static Timing Only

- Calibration delay at multiple voltage and temperature points
  - 4 chips tested
  - 4 paths per CPM, 5 CPMs per core, up to 8 cores per chip
- 13 delay step spread
- Heavy, steady-state workload
- Dark blue: 0.9V, 50C
- Magenta: 1.1V, 60C
- Cyan: 1.3V, 60C
- Yellow: 1.3V, 80C

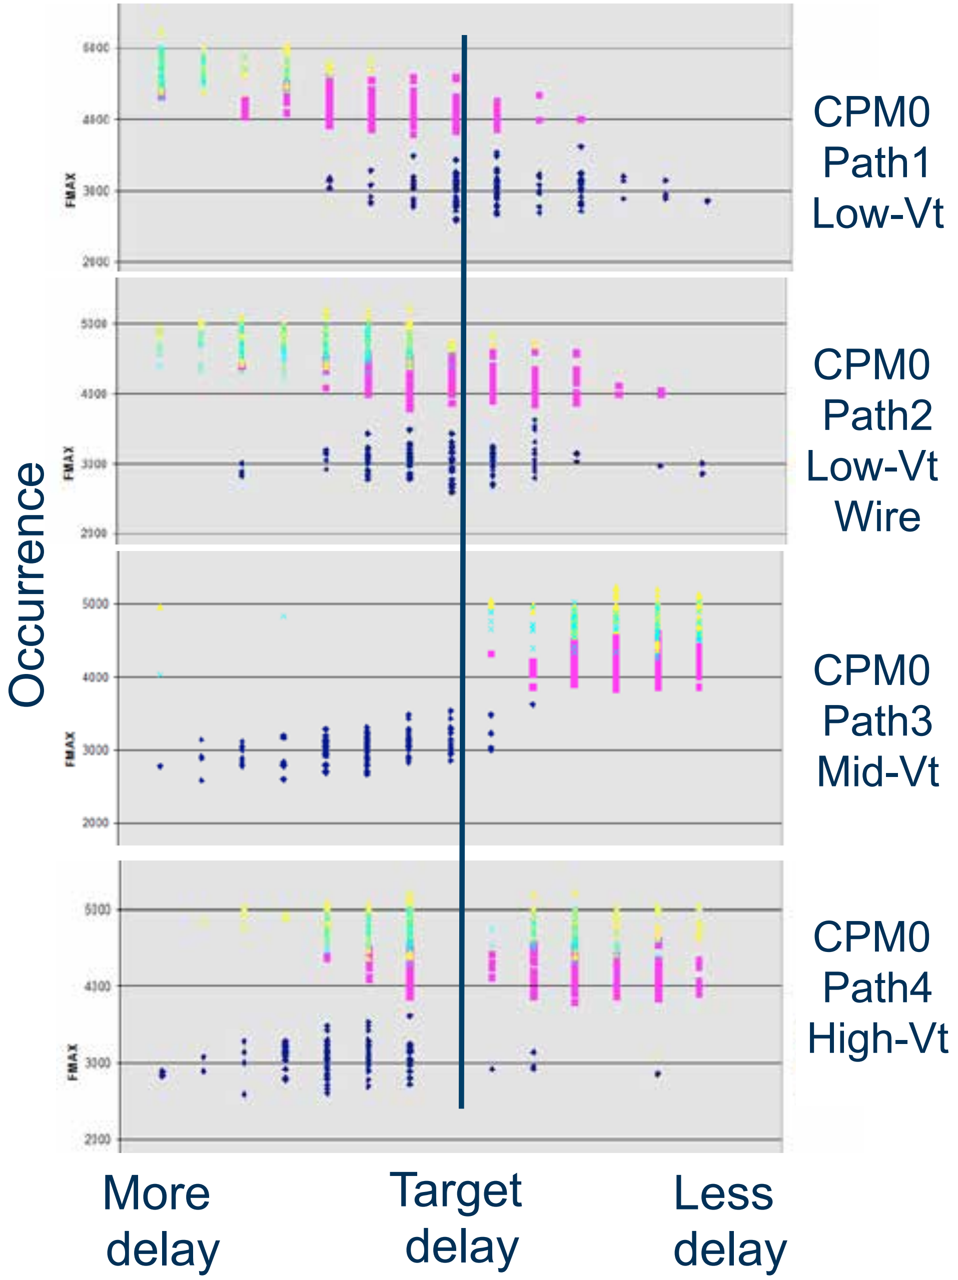

HW2: Static Timing Only After Adjusting for Path Offsets

- **Calibration delay at multiple voltages**
  - **More than 20 parts tested**
- **High-Vt distribution shifts right and lines up with Low-Vt inverter**
- **Low-Vt distributions shift left**
- **Mid-Vt distribution shifted right**

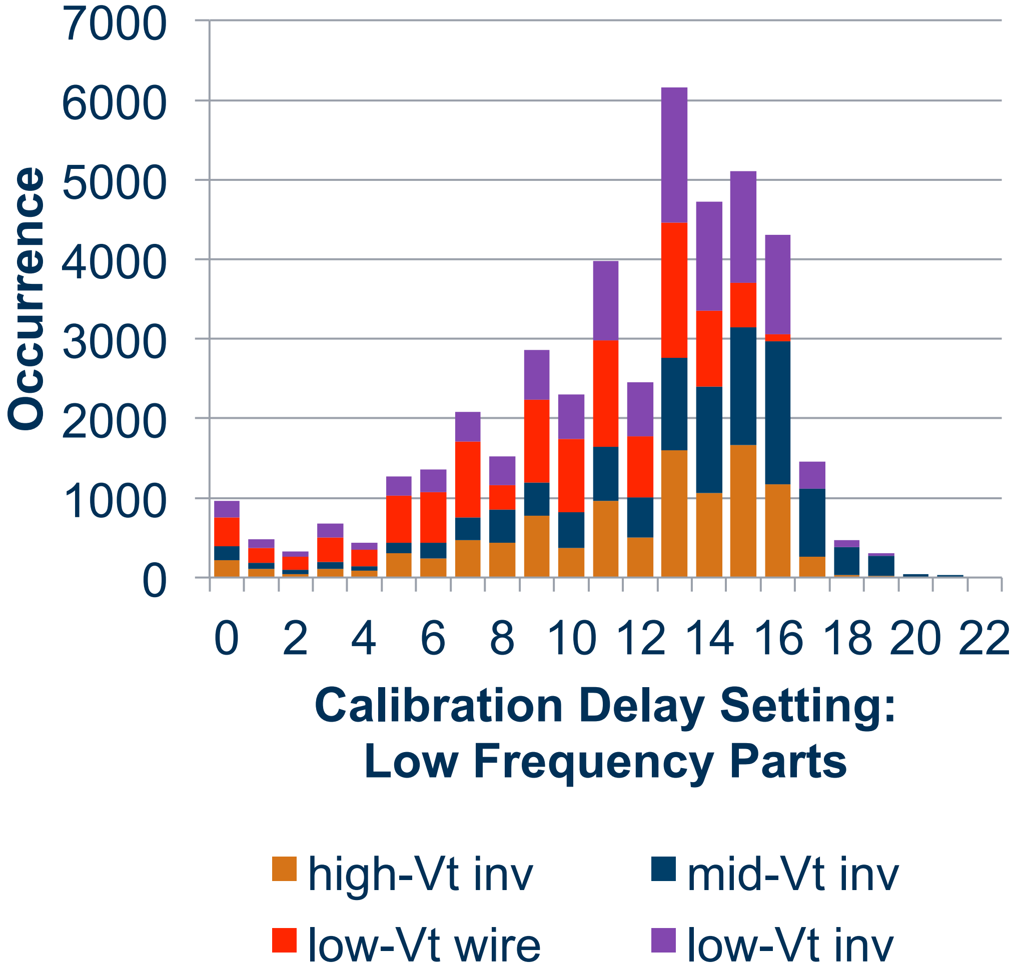

HW2: Bit-to-bit Calibration Delay Spreads All Operating Points

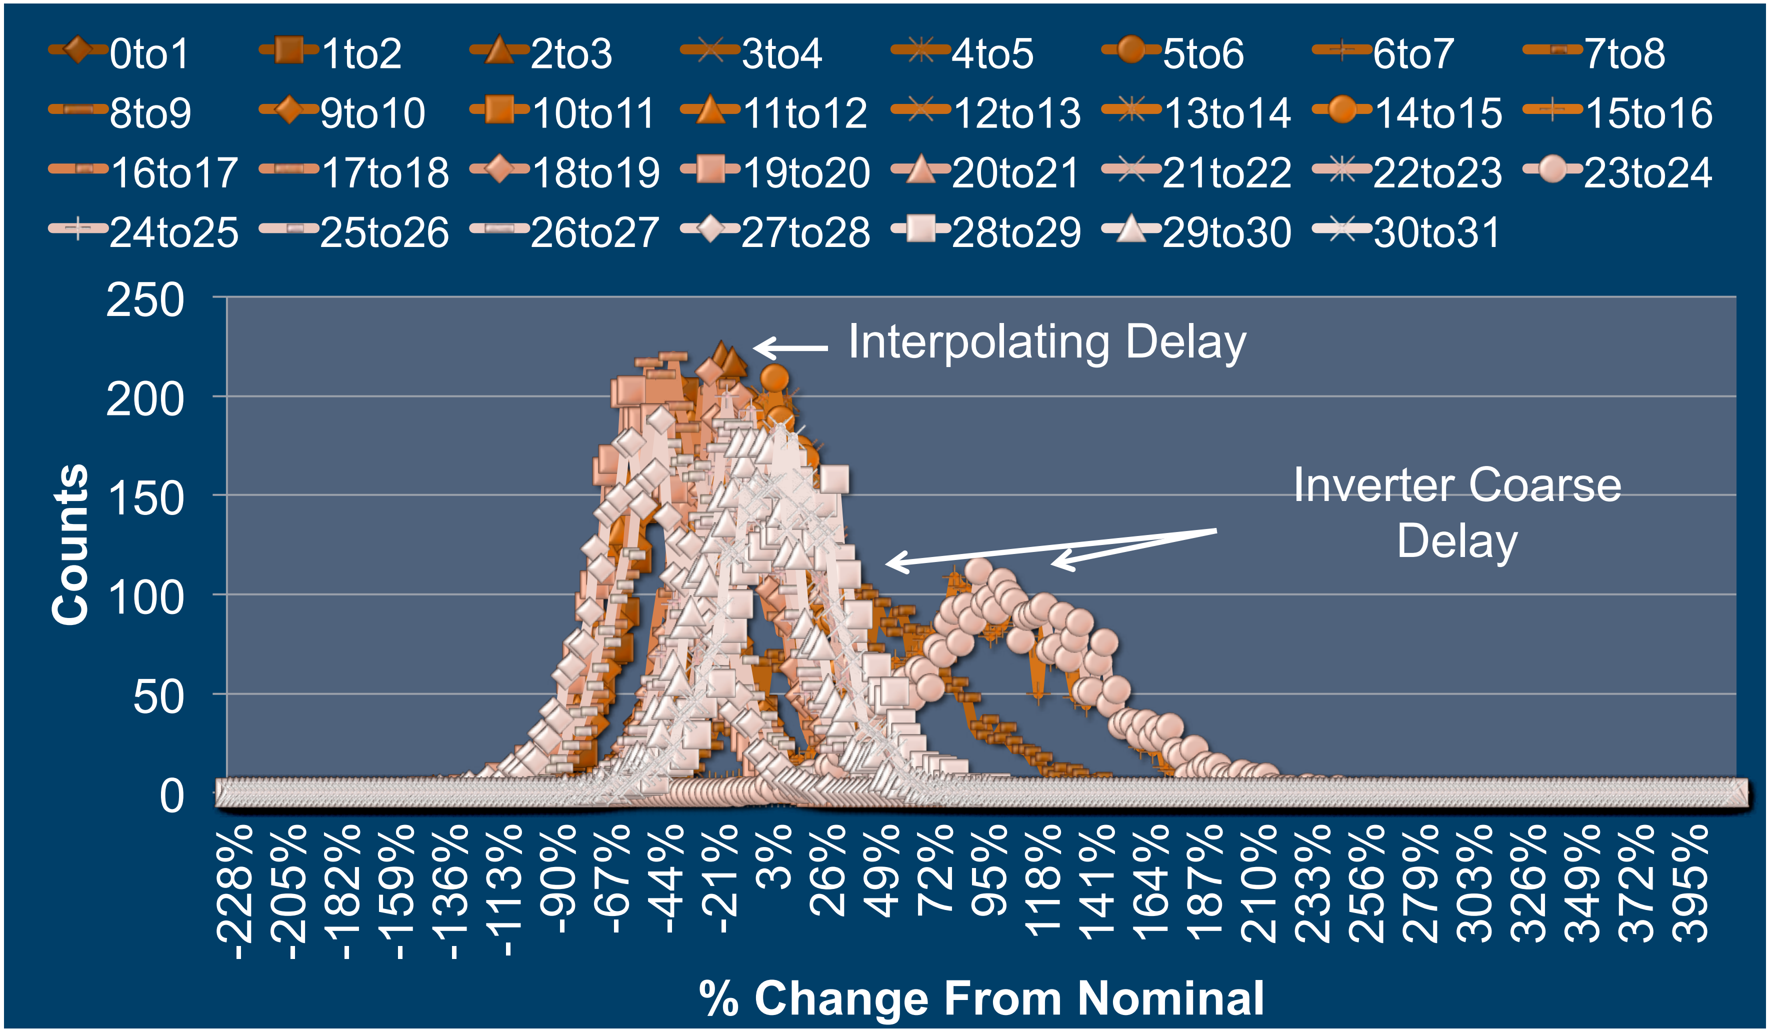

Measured μ and 3-σ of Calibration Delay Step Size

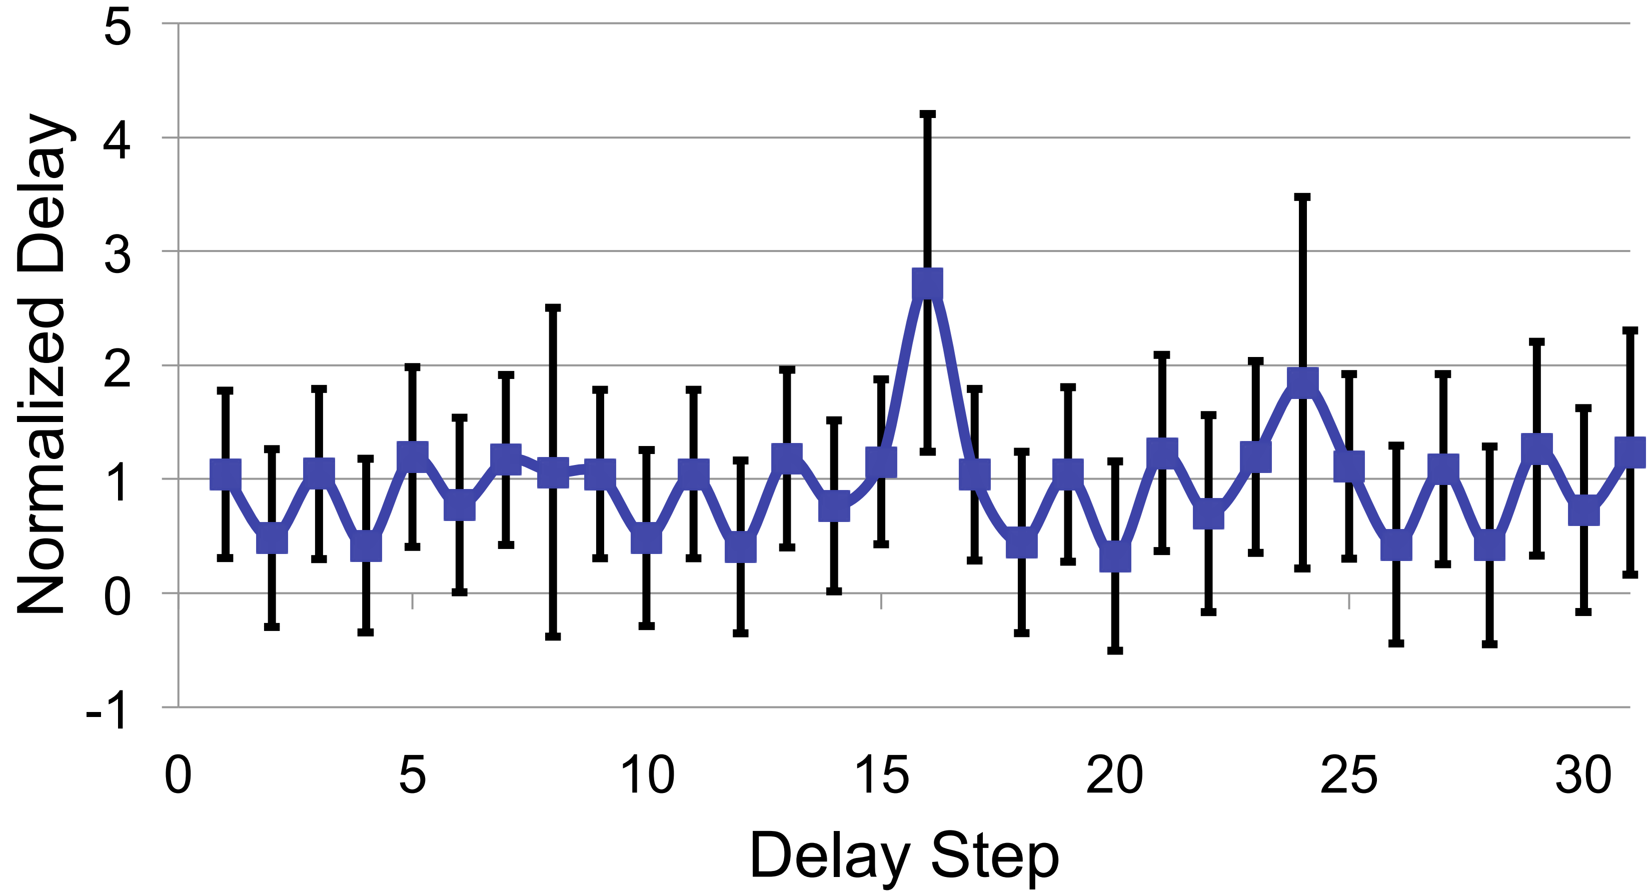

Calibration Range Breakdown

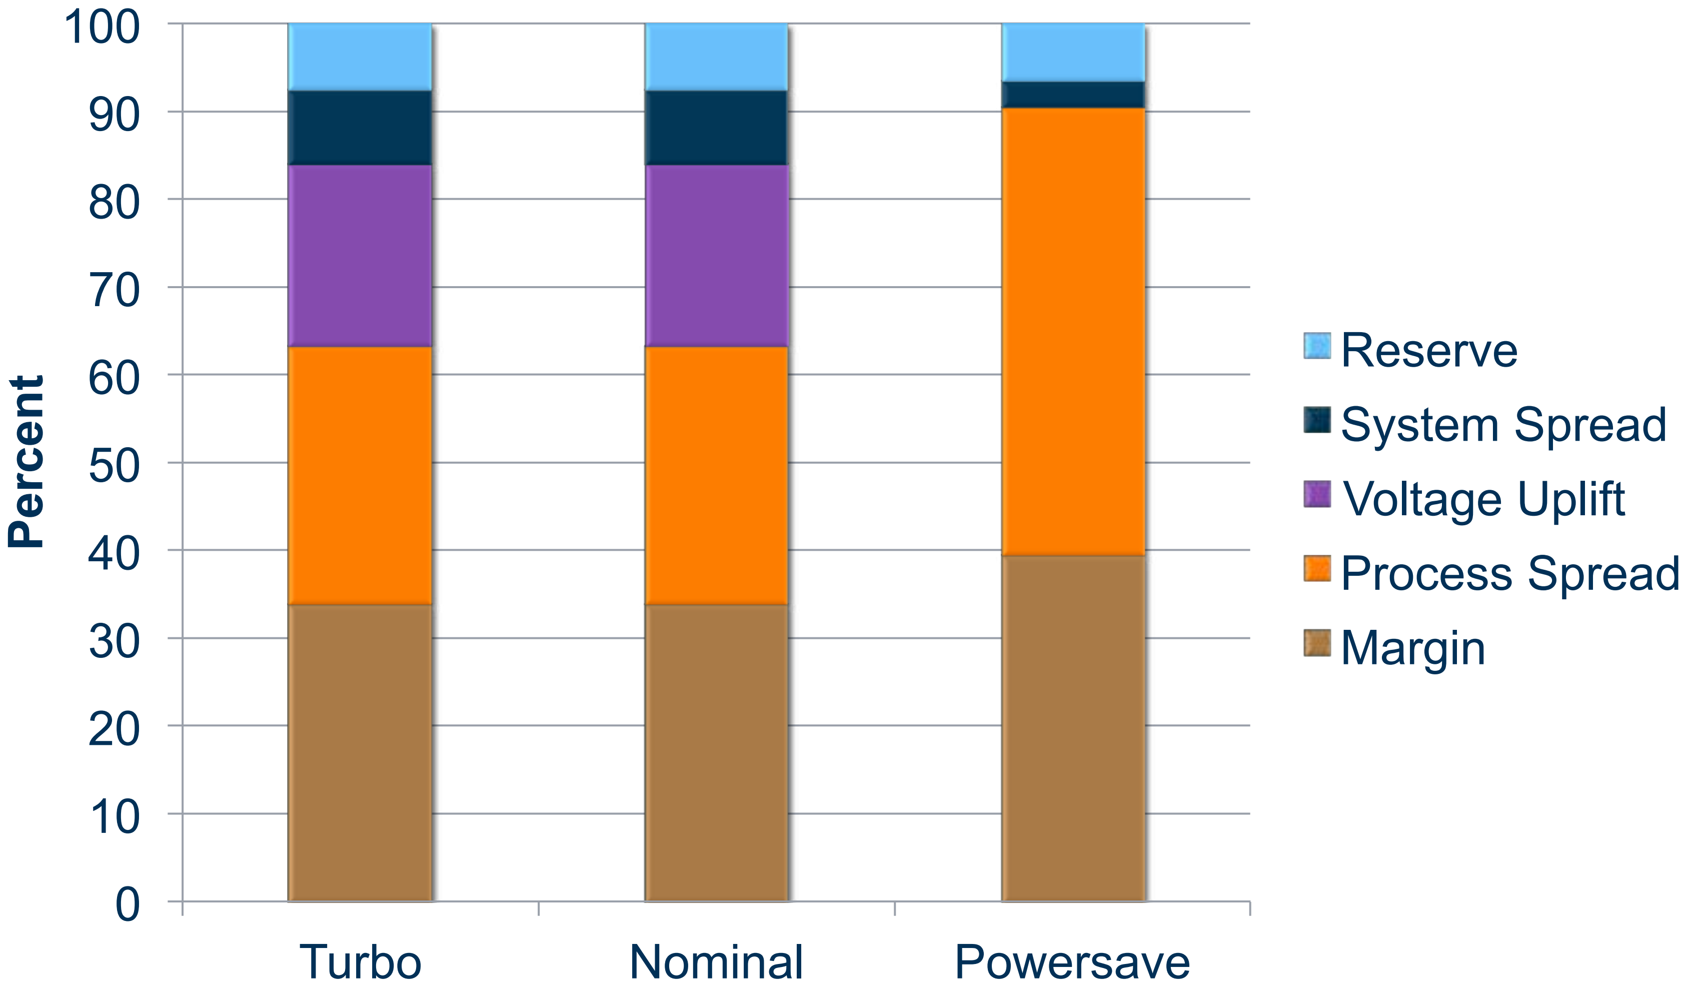

CPM Distribution

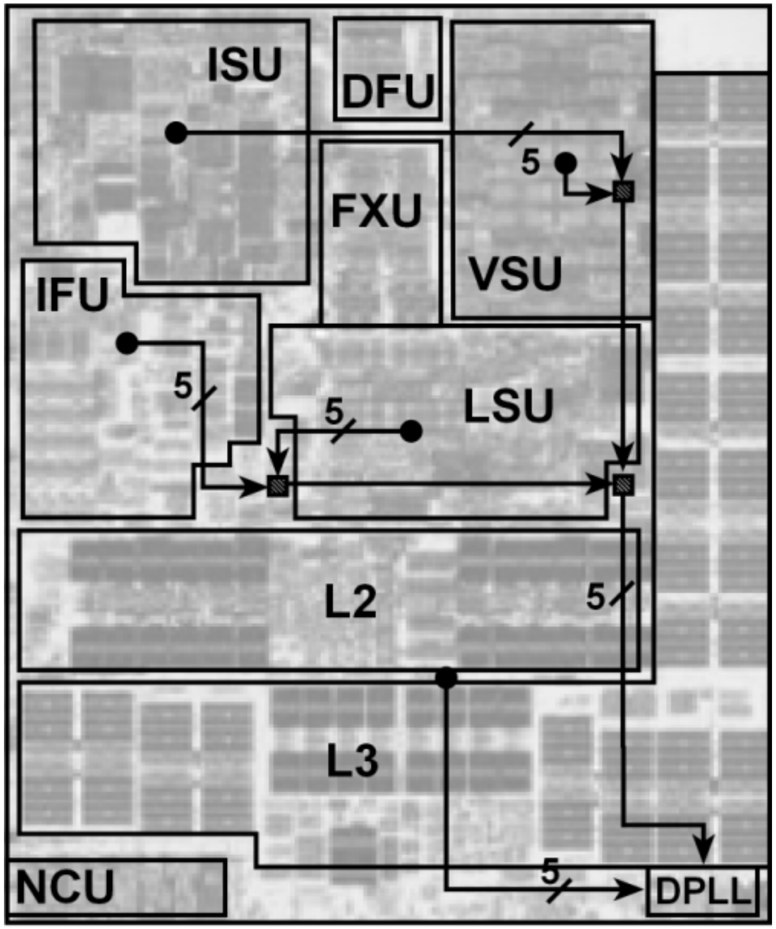

● = CPM (Critical Path Monitor)
▨ = AND buffer

- 4 Full-clock CPMs
- 1 Half-clock CPM
- Hierarchical
  - Control: 10724 transistors
  - Custom synthesis block: 5257 transistors
- Simulation
  - Spice
    - Slow (6hrs, 34mins)
  - Static timing
    - Fast (50 minutes)
    - Very accurate at 2 points: early and late mode timing

Measured μ and 3-σ of Delay Needed for Calibration

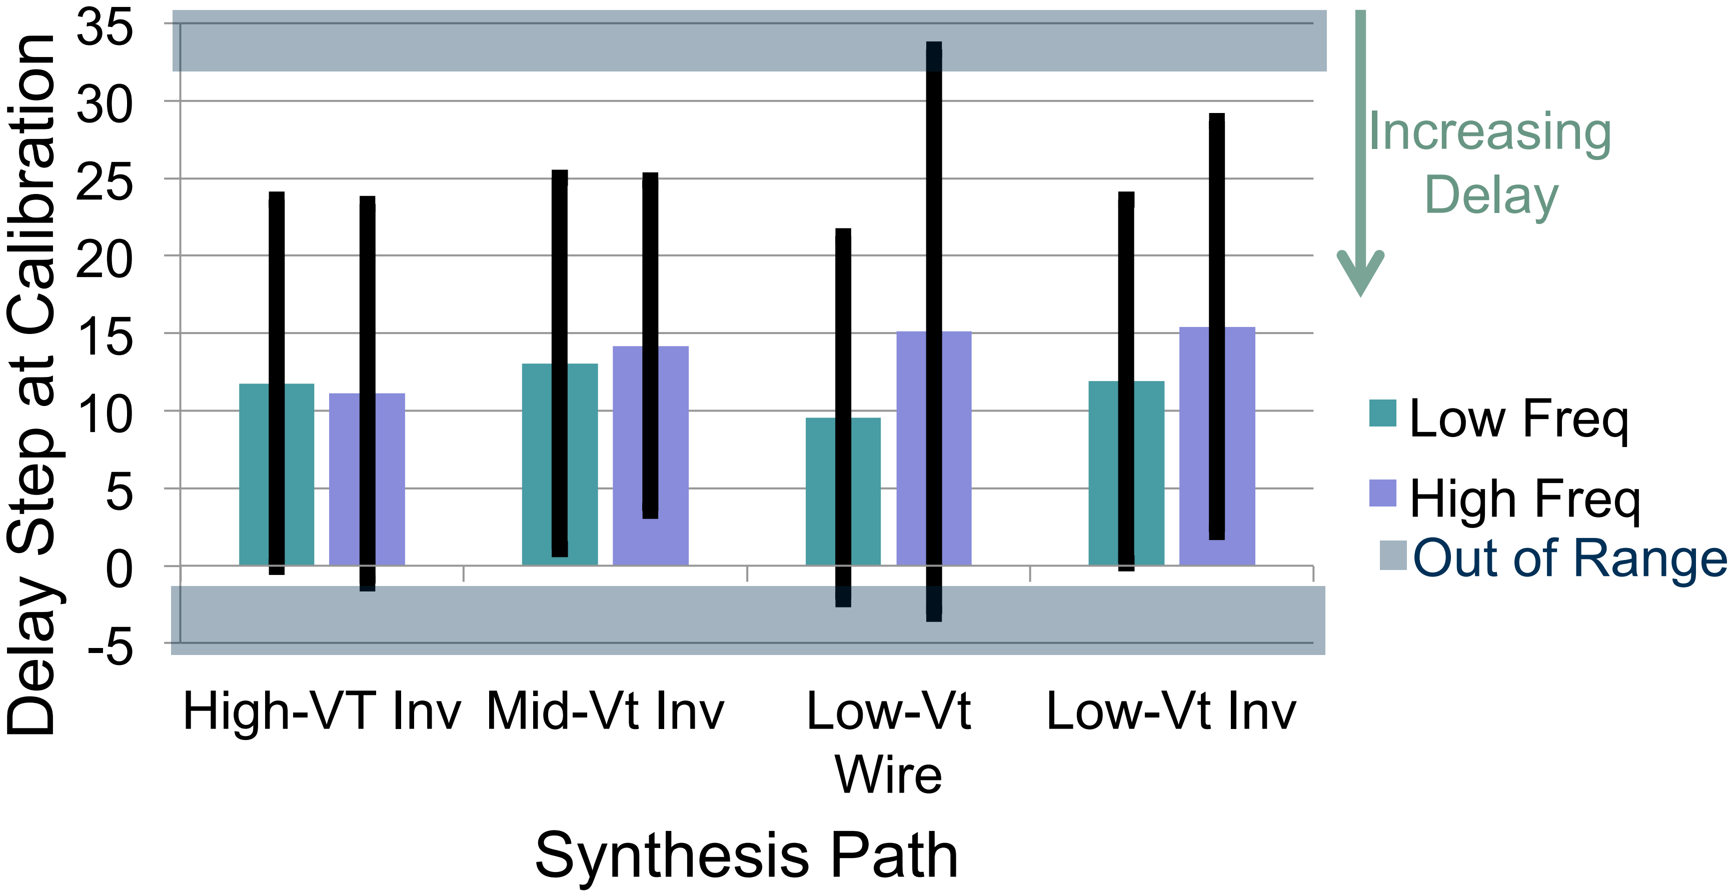

Measured Slopes vs. Frequency

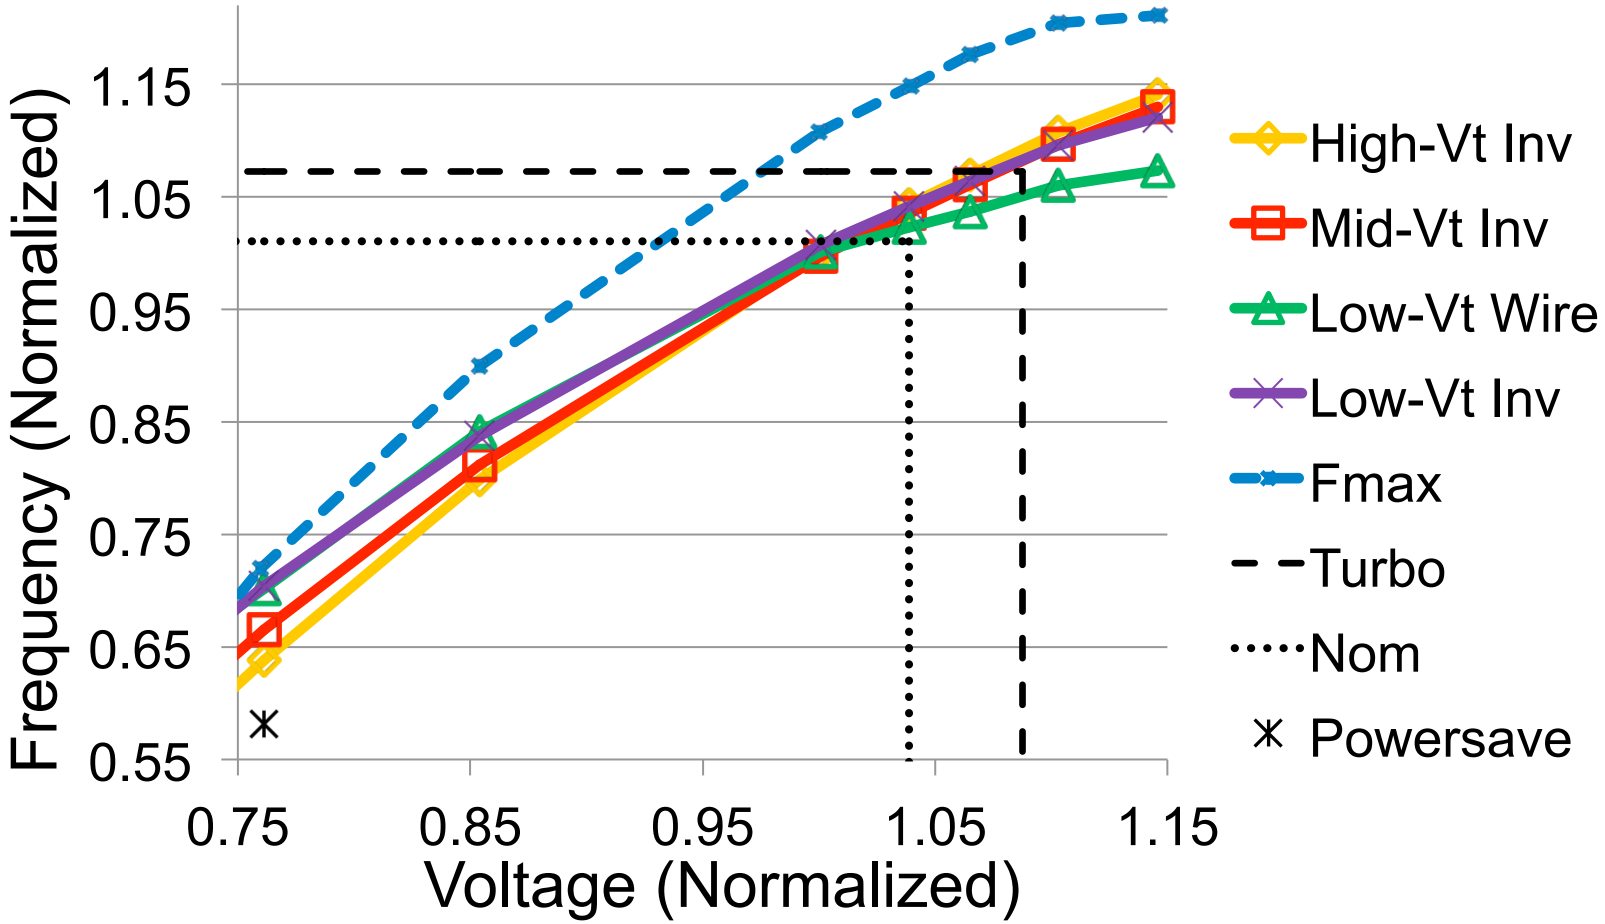

Resultant Deviation from Calibration Frequency

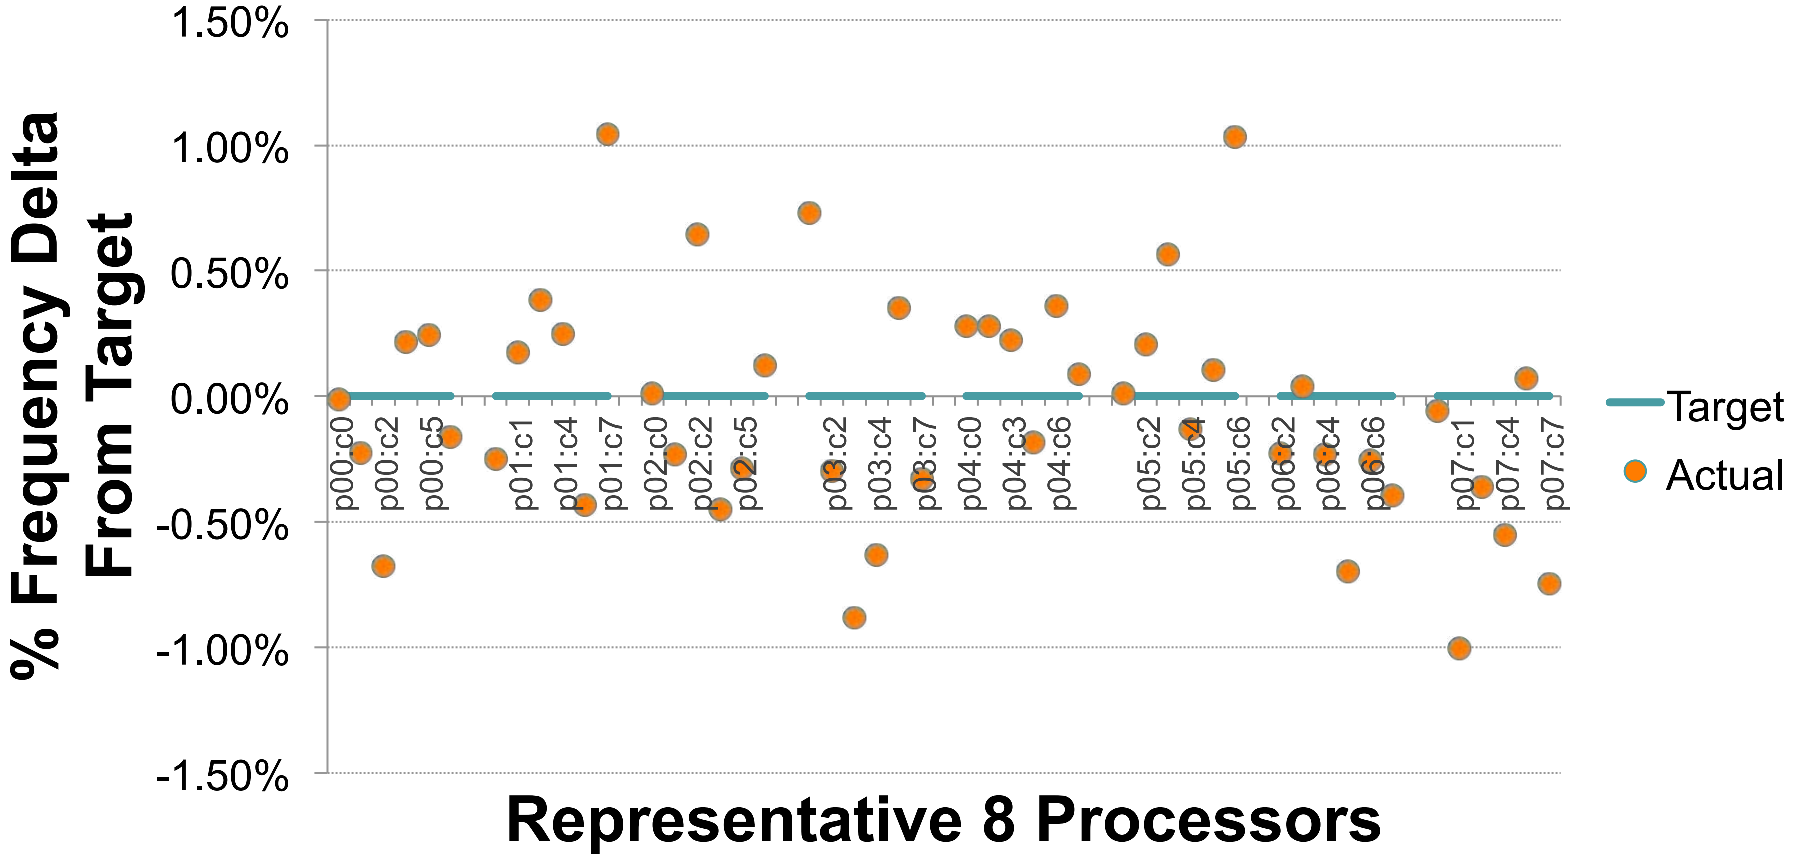

Simulated Range with More Detailed Extraction: Full Frequency Range

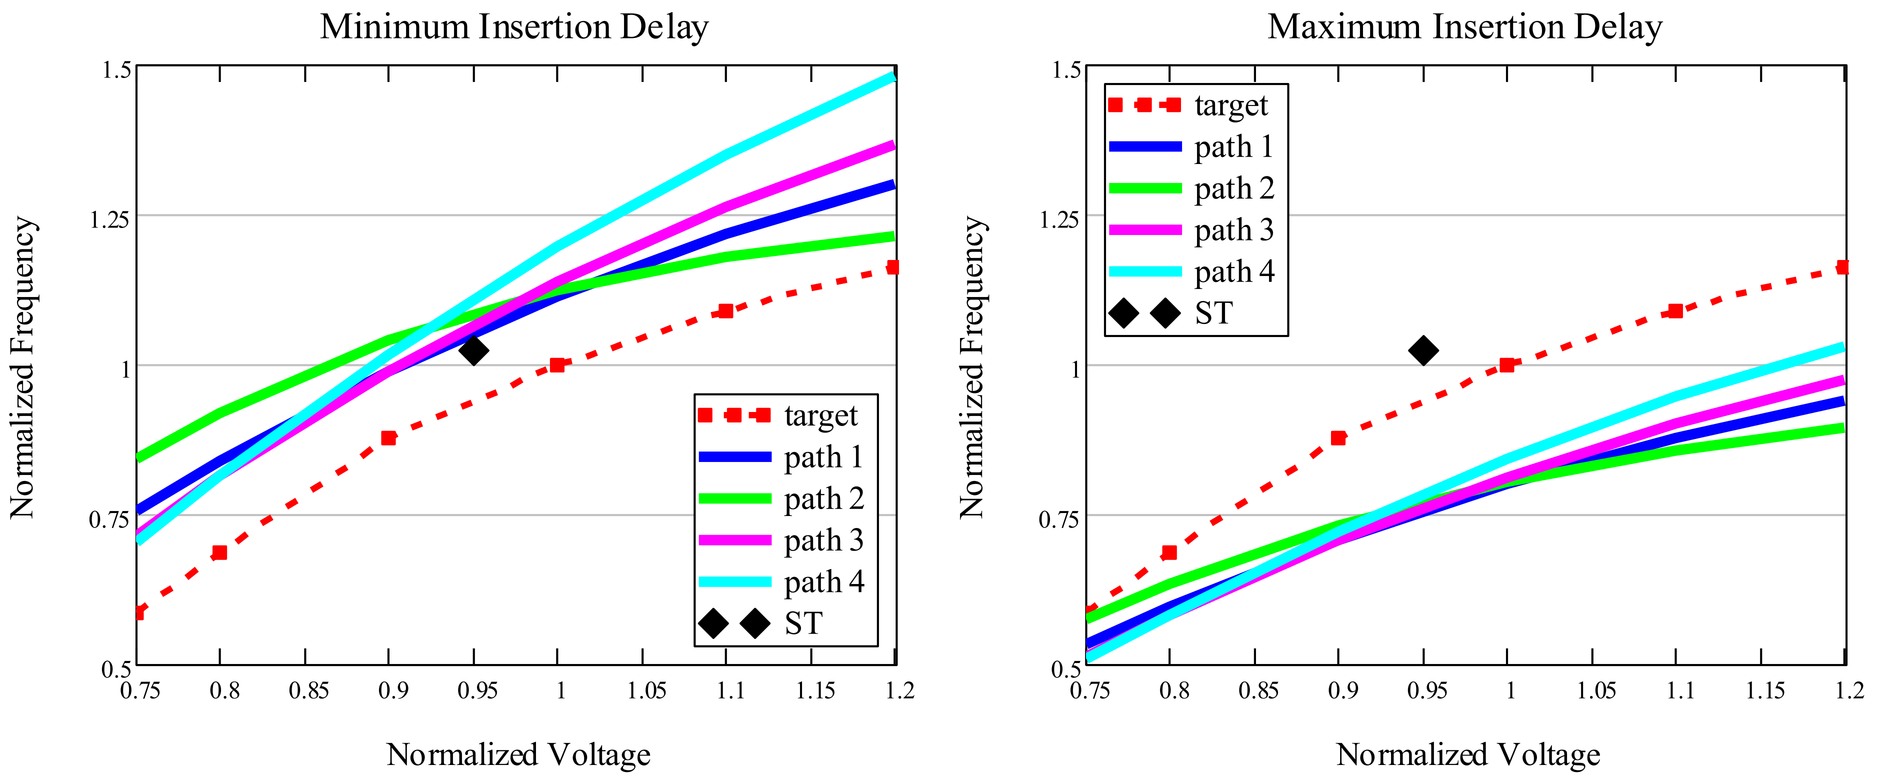

What can be done with circuits

- New capability to keep timing margin nearly constant
  - Convert excess timing margin into a voltage reduction
  - Reduce traditional voltage margin when conditions are not worst-case (Some voltage margin is retained for aging, calibration inaccuracy, etc.)

1. **Measure** excess operational margin with timing margin sensor
   - Difference from a calibrated reference point
2. **Protect** timing margin against voltage droop by adjusting frequency
   - Hardware-based **timing margin controller**
3. **Save energy** by converting excess timing margin into voltage reduction
   - Software-based **performance controller**

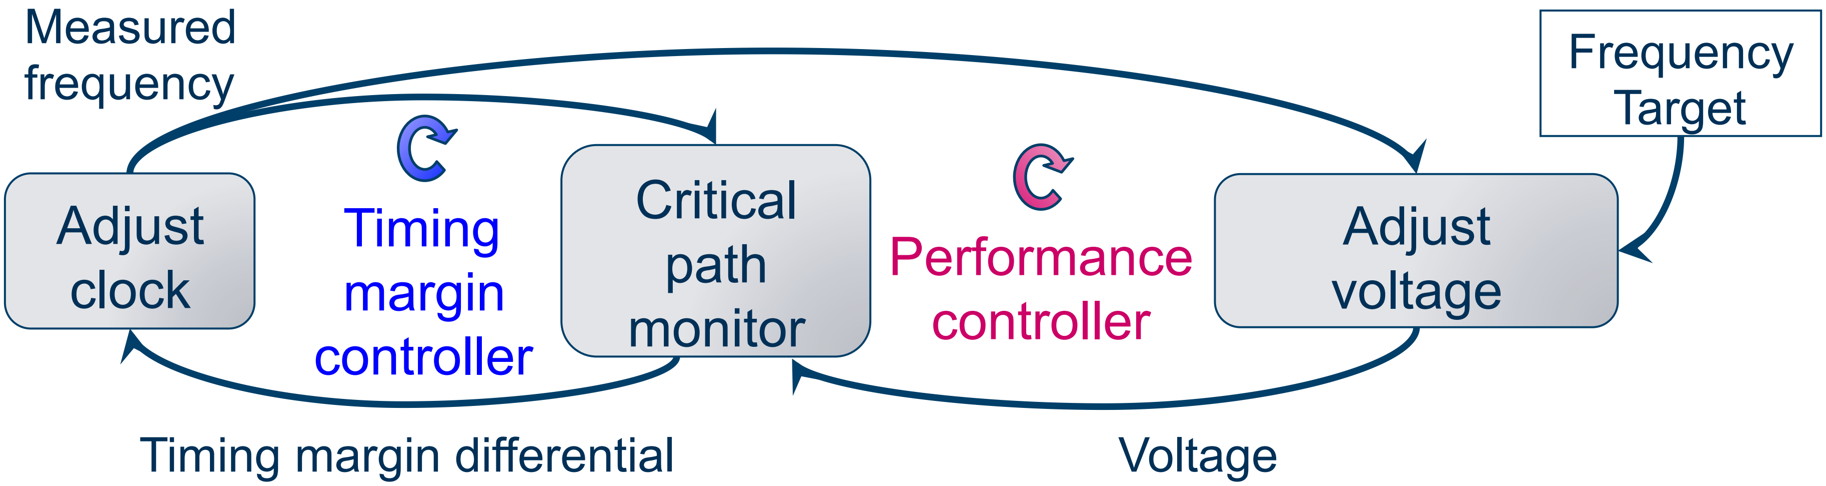

Measured frequency

Frequency Target

Adjust clock

Timing margin controller

Critical path monitor

Performance controller

Adjust voltage

Timing margin differential

Voltage

Use of On-die process-voltage-temperature sensors: long-term effects of process variation and aging on performance and reliability

Opportunity exists to overcome systematic variability through advanced modeling, sensors and monitors driven adaptation to improve efficiency.

- Real-time sensors
- Advanced real-time power management control systems
- Calibration with wafer level process data with characterizations structures

Real-time process, voltage, temperature and timing sensors

EnergyScale management system for P-series

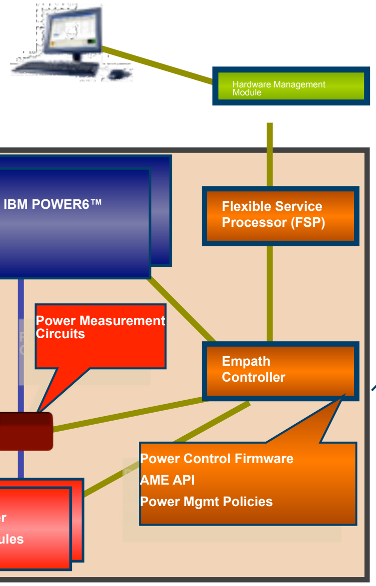

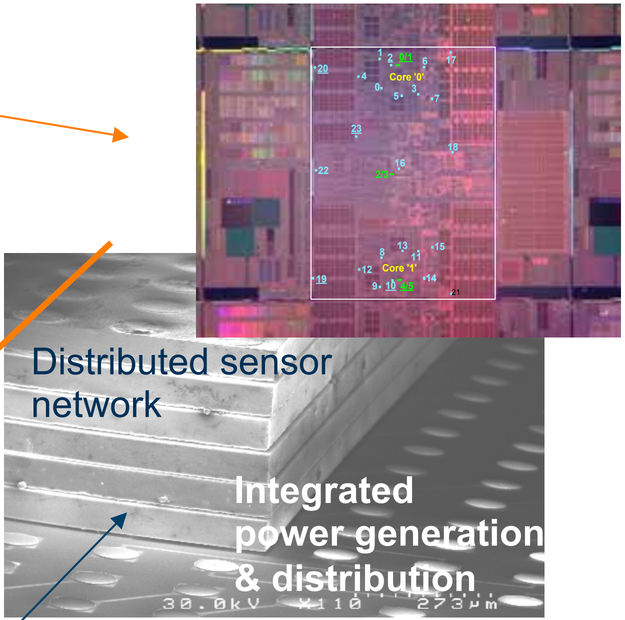

Distributed sensor network

Integrated power generation & distribution

Process variability sensors

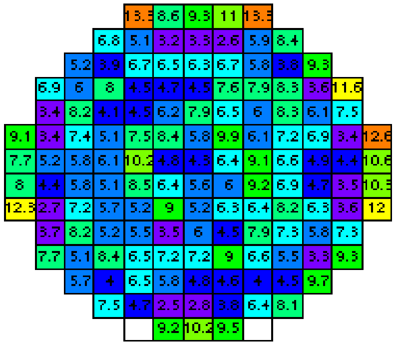

				13.3	8.6	9.3	11	13.3				
			6.8	5.1	3.2	3.3	2.6	5.9	8.4			
		5.2	3.9	6.7	6.5	6.3	6.7	5.8	3.8	9.3		
	6.9	6	8	4.5	4.7	4.5	7.6	7.9	8.3	3.6	11.6	
	3.4	8.2	4.1	4.5	6.2	7.9	6.5	6	8.3	6.1	7.5	
9.1	3.4	7.4	5.1	7.5	8.4	5.8	9.9	6.1	7.2	6.9	3.4	12.6
7.7	5.2	5.8	6.1	10.2	4.8	4.3	6.4	9.1	6.6	4.9	4.4	10.6
8	4.4	5.8	5.1	8.5	6.4	5.6	6	9.2	6.9	4.7	3.5	10.3
12.3	2.7	7.2	5.7	5.2	9	5.2	6.3	6.4	8.2	6.3	3.6	12
	3.7	8.2	5.2	5.5	3.5	6	4.5	7.9	7.3	5.8	7.3	
	7.7	5.1	8.4	6.5	7.2	7.2	9	6.6	5.5	3.3	9.3	
		5.7	4	6.5	5.8	4.8	4.6	4	4.5	9.7		
			7.5	4.7	2.5	2.8	3.8	6.4	8.1			
					9.2	10.2	9.5					

Live correlated field data for comparison to historical process data

Design Example:

- Specifications
  - 1A/core
  - 500mA/ns
  - 1V supply
  - 12ps FO1 inverter delay
  - 100pF on-chip de-coupling
  - 4GHz clock
  - 5 cycle response time from circuit to clock change
  - 1nH inductance to regulator
  - 10mOhm series resistance in supply grid