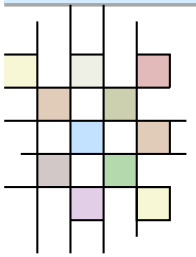

Dependability Issues Due to Scaling Towards Nanometer Size Devices:

Aggressive and Adaptive Mitigation Techniques Maybe Key for Solution

Arun K. Somani

Dependable Computing and Networking Laboratory
Department of Electrical and Computer Engineering
Iowa State University, Ames, IA, 50011

arun@iastate.edu

IOWA STATE UNIVERSITY

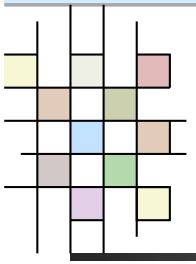

Technology Scaling

- Every 30% downscaling of technology node
  - Transistor density doubles
  - Gate delay reduces 30%
  - Operating frequency improves 43%
  - Active power consumption halves
  - 65% energy savings
- Frequency scaling inhibited with recent generations
  - Low power requirements
  - Process variations
  - Reliability concerns
- High speed, low leakage requirements
  - Determines the choice of supply and threshold voltages

How the Progress is Holding Up?

- Drives semiconductor performance
- Enables newer technologies

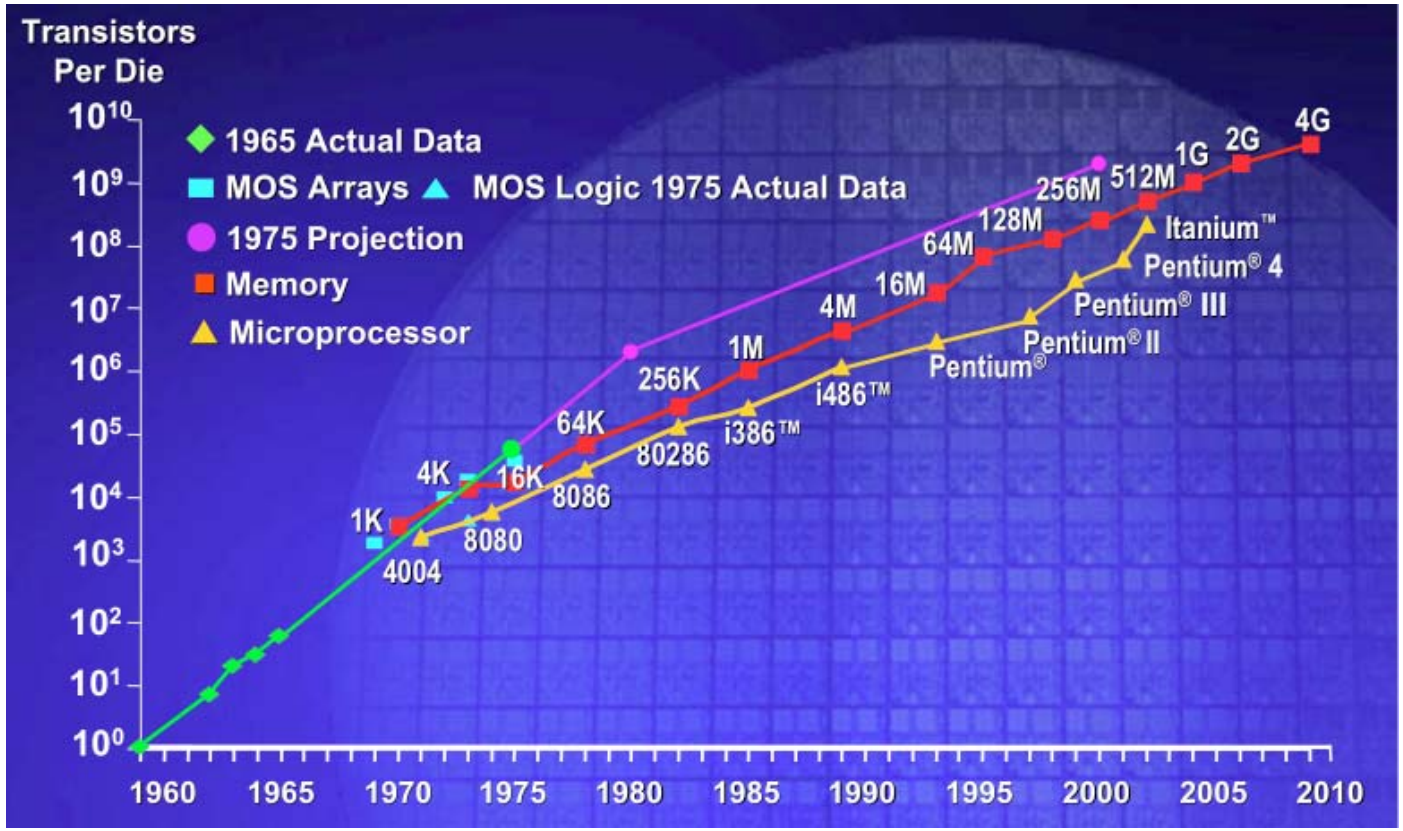

Source: Intel

A Few Things Are Here to Stay

- Leakage Power in MOSFETs
  - Sufficient overdrive required for high speed switching
  - Lower V_T leads to more leakage
- Gate Leakage
  - Tunneling current through gate dielectric
  - High-k dielectrics used in 45nm technology
    - Arrest gate leakage
- Process variations increase with scaling
  - Random and systematic variations in delay, power, yield
  - $V_t \downarrow \rightarrow$ Delay $\downarrow$, $L_{eff} \uparrow \rightarrow$ Delay $\uparrow$, $V_{dd} \downarrow \rightarrow$ Delay $\uparrow$, T $\uparrow \rightarrow$ Delay $\uparrow$
- Thermal Variation

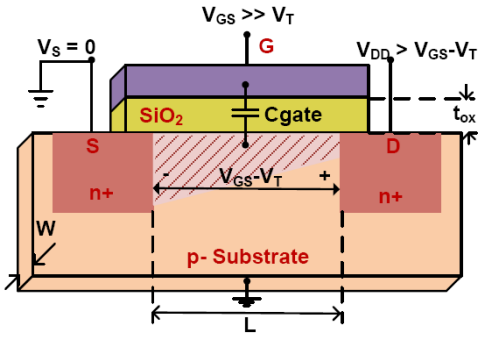

Temperature Variations

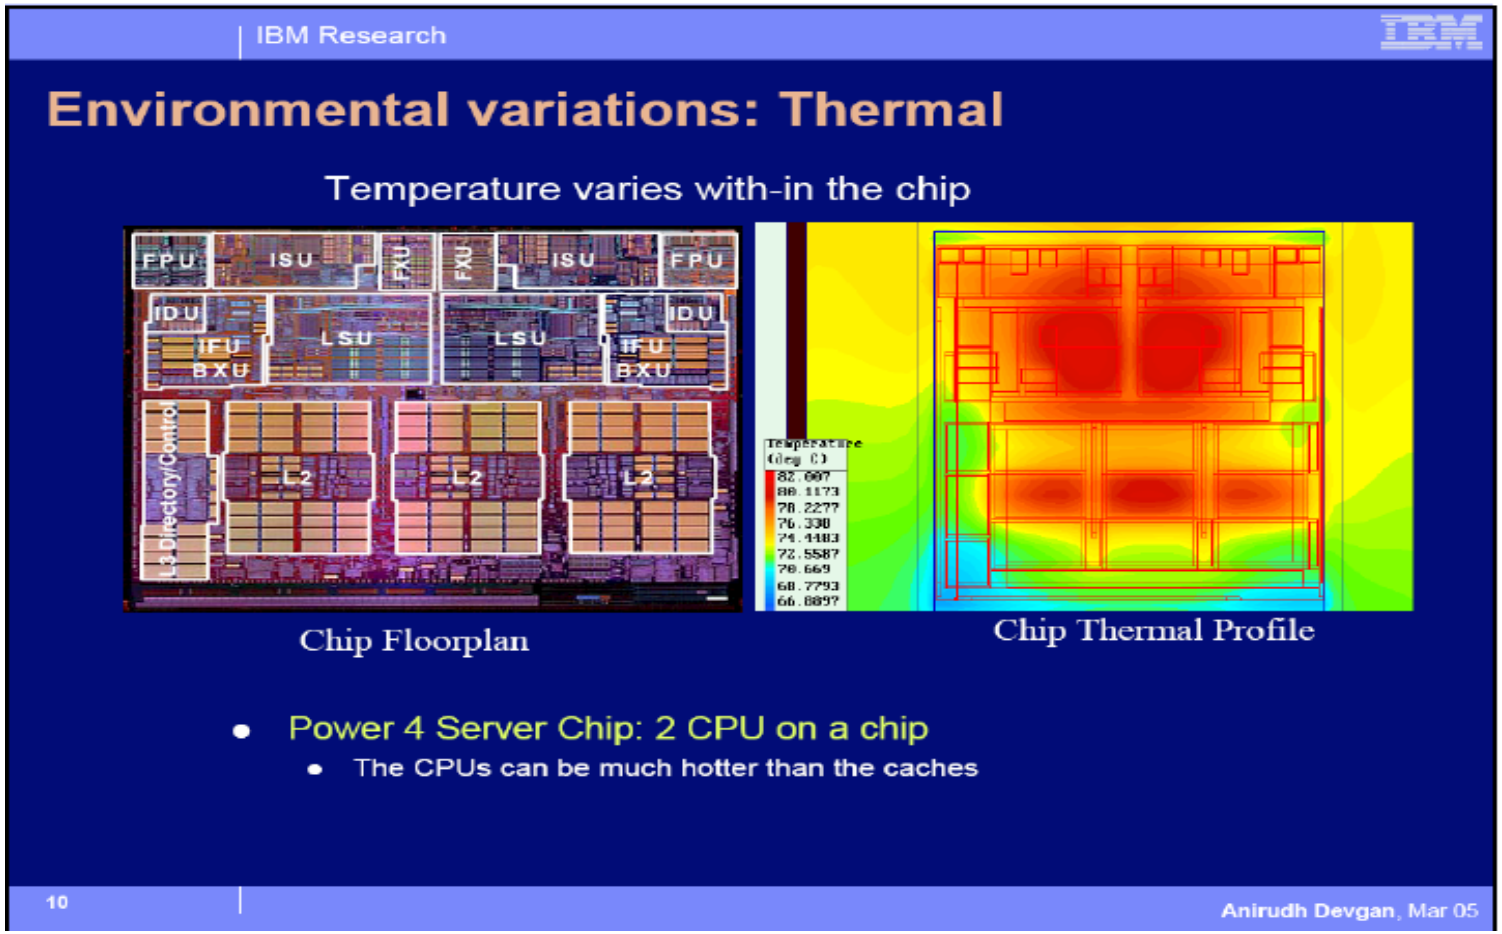

IOWA STATE UNIVERSITY

Original Source: Anirudh Devgan, IBM Research

Challenges for Future Manufacturing

- Ultimate limit 0.3 nm (Silicon atoms distance)
  - Various barriers seen over time
  - Overcome with changes in material and process technology
- Degradation of performance with downscaling
  - Interconnect delay increases with increase in resistance and capacitance of narrow and dense metal lines
- Higher power consumption will continue as a problem
- Unaffordable manufacturing cost for smaller sizes
  - Semiconductor companies moving towards fab-lite model
  - Yield and the time-to-market with newer technologies is becoming longer

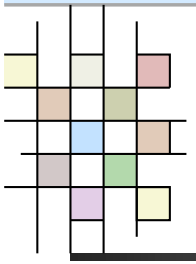

What to Look Forward For?

- Error tolerance rather than avoidance
- Built in fault tolerance for all designs
- Selective replication instead of full scale redundancy
- Design adaptability
  - Key for low overhead solutions
- Design optimizations
  - Dynamic schemes
  - Possible through speculation

Reliable Overclocking (Aggressive Designs)

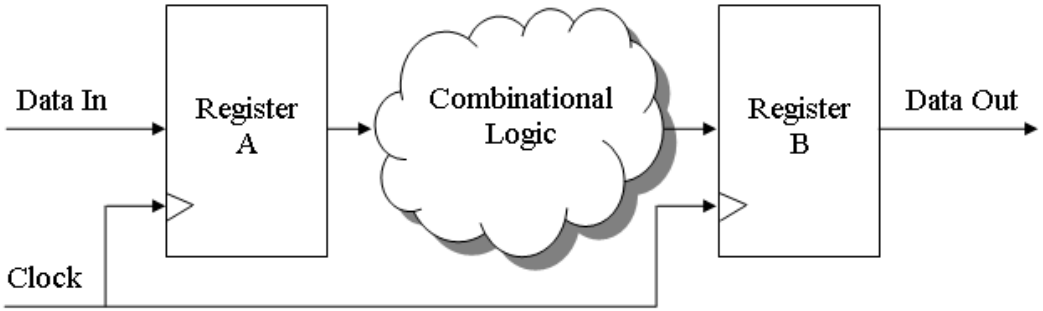

- Typically clock period is determined by the maximum delay from A to B which depends physical implementation, operating environment, and temperature and supply voltage variations
- Traditionally, worst case delays assumed
  - Result - overly conservative clock period
- Pipelined processor
  - Longest/slowest stage limits the period of the entire pipeline

Reliable Overclocking (Aggressive Designs) – Contd.

- Problem to address in nanometer design space
  - Provide high performance by exploiting PVT variations
  - Enhance system dependability with low cost solutions
- Clock beyond worst case delay, relying on data dependent delays
- Timing errors may occur at overclocked speeds
- Aggressive, but reliable, design methodologies employ relevant timing error detection and recovery schemes
  - Razor-Micro'03, Sprite-DSN'07
- Performance 15-20%, Error rate below 1%
- Safety critical systems, real-time constraints supported

Why Past Solutions are not Acceptable

- Traditional techniques
  - TMR solutions incur high cost and performance penalty
  - Dual latching dynamic optimization uses less area
  - False positives and high penalty for error recovery are concerns
- Static power Vs Dynamic power
  - Both are comparable for today's technology
  - Thus logic replication is not a viable alternative

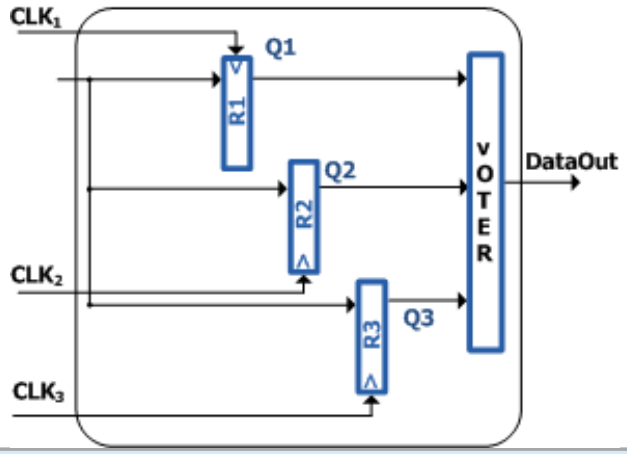

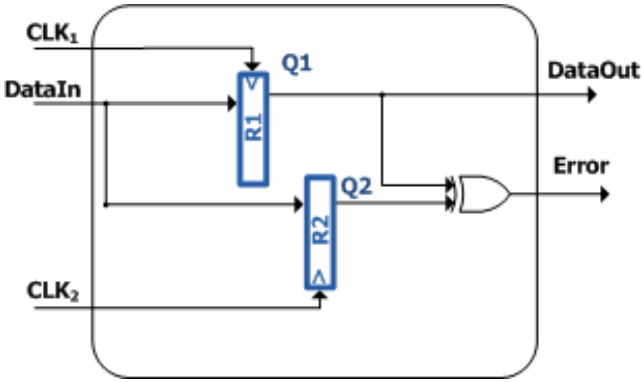

Offering More Design Features with Added Redundancy

- Soft Error Mitigation, SEM [DSN'09]
  - Circuit level speculation, local recovery, no false positives, high fault coverage (like TMR tolerates both SEU and SET)
  - No performance overhead, operating frequency $f_{sys} \leq 1/t_{pd}$
- Soft and Timing Error Mitigation, STEM [DSN'09]
  - Like SEM, but detects and correct timing errors
  - Can be deployed in aggressive system designs
  - Timing speculation, like overclocking [DSN'07] and DVS [MICRO'03]

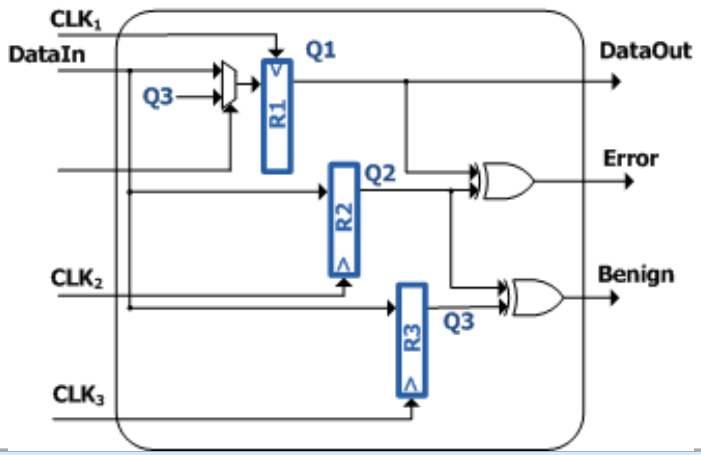

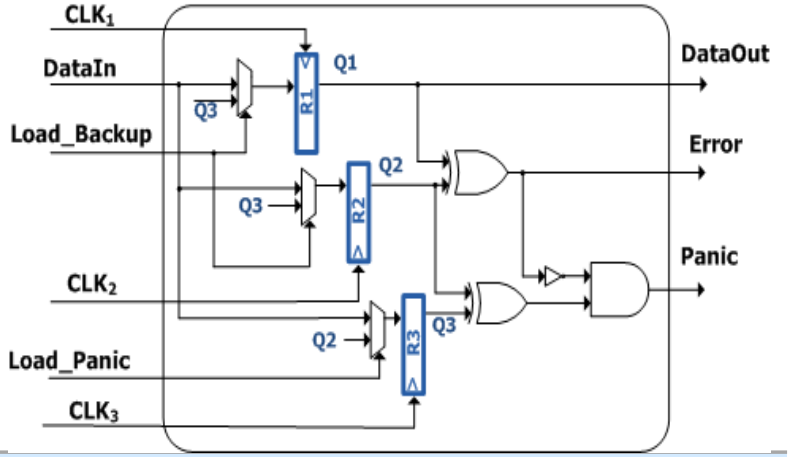

Design Constraints

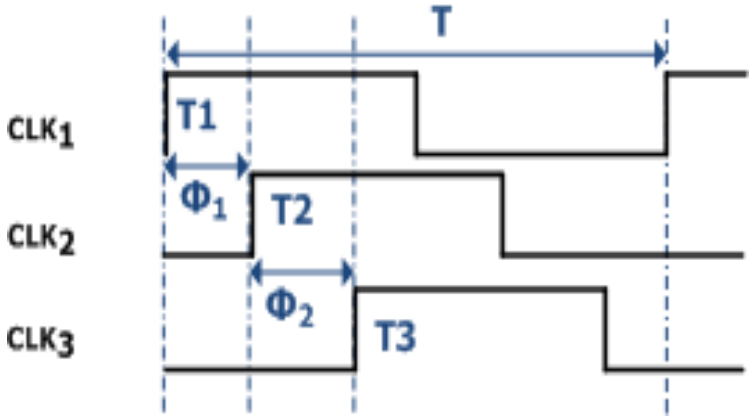

$$\Phi_1 = T_2 - T_1 \geq T_{PW} \quad (5)$$

$$\Phi_2 = T_3 - T_2 \geq T_{PW} \quad (6)$$

$$T_{CD} \geq \Phi_1 + \Phi_2 \quad (7)$$

$$T + \Phi_1 \geq T_{PD} \quad (8)$$

T_{CD} = Contamination delay of the logic circuit

T_{PD} = Propagation delay of the logic circuit

T_{PW} = Expected soft error/noise pulse width

Φ_1 = Phase shift between CLK_1 and CLK_2

Φ_2 = Phase shift between CLK_2 and CLK_3

T = Clock period

Dynamic Frequency Scaling

- Clock frequency is scaled while satisfying the error rate constraint
- Limits of DFS
  - F_{MAX} (Minimum possible frequency)
    - Set by worst-case design settings
  - F_{MIN} (Maximum possible frequency)
    - As shown in equation (11)

$$T_{CD} \geq D_2 \quad (9)$$

$$D_2 - D_1 \geq T_{PW} \quad (10)$$

$$T_{MIN} + D_1 \geq T_{PD} \quad (11)$$

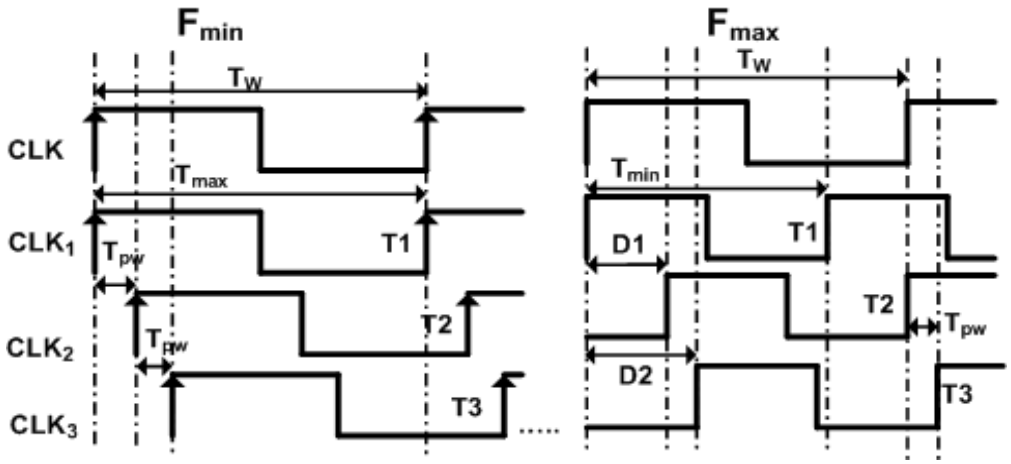

T_{CD} = Contamination delay of the logic circuit

T_{PD} = Propagation delay of the logic circuit

T_{PW} = Expected soft error/noise pulse width

D_1 = Phase shift between CLK_1 and CLK_2

D_2 = Phase shift between CLK_2 and CLK_3

Pipeline Design

- Using STEM
  - Input clocks are constrained to provide fault tolerance
  - Extra buffer stage to ensure only “gold” data to memory
- Stage error signal: Generated from error signal in that stage
- Global error signal is generated from all stages
- Error rates are monitored and used by clock unit

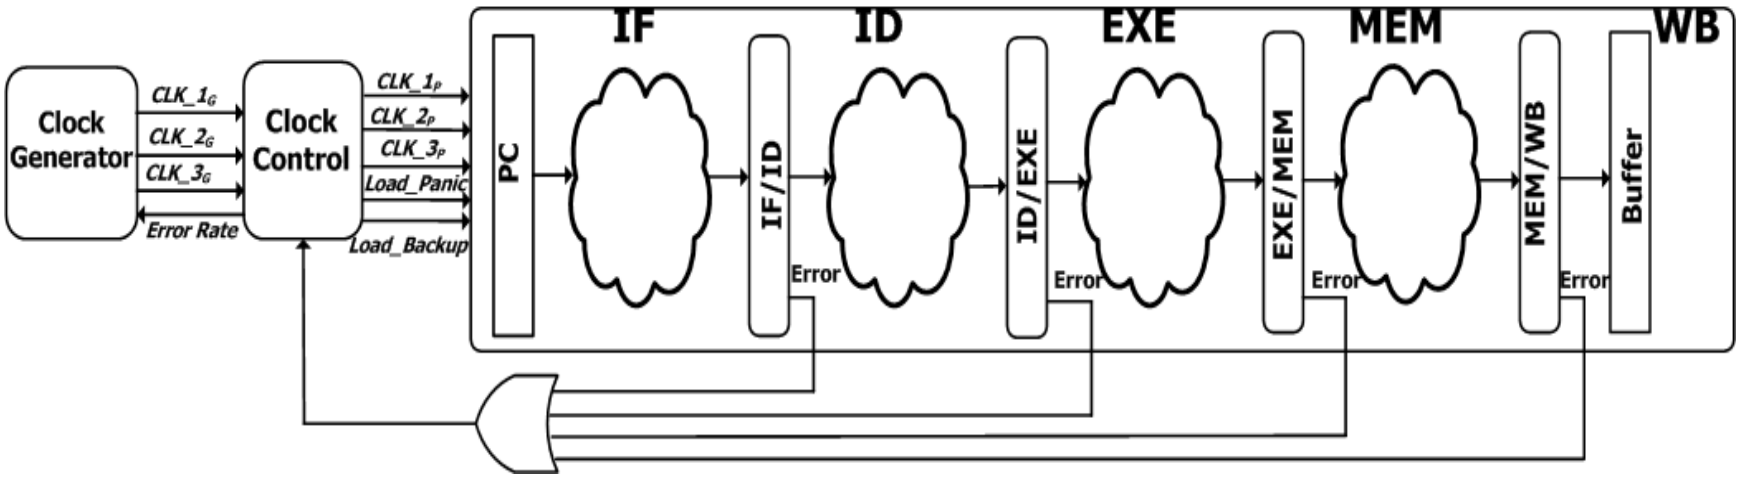

Performance Analysis

- Limiting factor for frequency scaling
  - With frequency scaling, no. of input combinations resulting in greater delays than the new clock period increases

$$N \times t_{ov} + n \times N \times k \times t_{ov} < N \times t_{wc}$$

$$k < (t_{wc} - t_{ov}) / (n \times t_{ov})$$

Notation:

t_{wc} : worst case clock period
t_{ov} : overclocked clock period
n : no of cycles to recover
N : total cycles required
k : error rate

- For STEM cells
  - 15% increase in frequency, error rate needs to be > 5.76% to yield no performance improvement
  - For error rates < 1%, a 2.6% increase in frequency is required to compensate the penalty paid for error correction

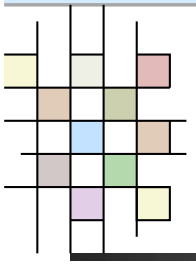

Three Interdependent Concerns

- Performance
  - Device scaling
  - Architectural innovations
  - Better-than-worst-case designs
- Dependability
  - Soft errors, silicon defects
  - Fault mitigation techniques
- Power Consumption
  - Low power design
  - Adaptive control mechanisms
- All managed through aggressive design methodology