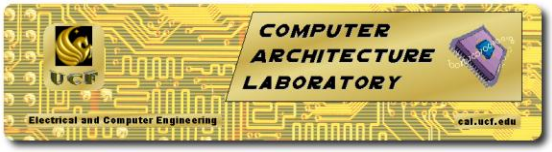

Extending Modular Redundancy to NTV: Costs and Limits of Resiliency at Reduced Supply Voltage

Rizwan A. Ashraf, A. Al-Zahrani, and Ronald F. DeMara

Department of Electrical Engineering and Computer Science
University of Central Florida
Orlando, FL

WNTC 2014 - 14 June 2014

Agenda

- Pros and Cons of ***Near-Threshold Computing (NTC)***
- Towards the Goal of Simultaneous Increase in ***Resilience*** and ***Energy Efficiency***
- Impact of ***Performance Variability*** for *N*-MR Systems
- Experimental Setup
- ***Energy Cost*** of Mitigating Variability
- Conclusions and Future Work

2

Increasing Interest in Near-Threshold Computing: *Limits*

- *Voltage Scaling* is a very effective way to reduce energy consumption
  - Total Energy = $E_{dynamic}$ + E_{static}
  - Dynamic Energy directly proportional to V_{DD}^2
  - E_{static} proportional to (V_{DD}* Leakage current * T_{clk})
- Extreme reduction of V_{DD}
  - Sub-threshold region
  - Theoretical Lower Limit of V_{DD} is 36 mV [1]
  - >12X Energy Savings as compared to Nominal
  - Massive Performance Penalty (exponential)
  - Limited Applicability

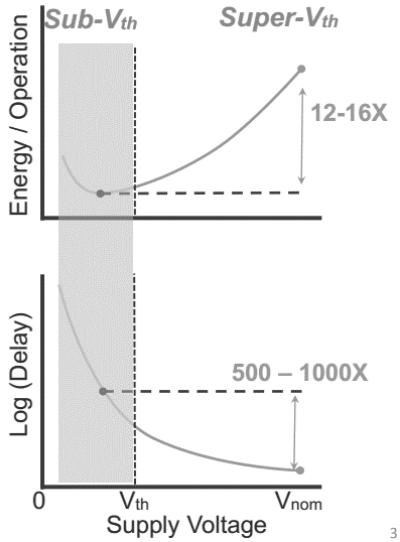

[Fick 2012] http://web.eecs.umich.edu/~mfojtik/fick_isscc2012_slides.pdf

3

Increasing Interest in Near-Threshold Computing: *In-Practice*

- *Voltage Scaling* is a very effective way to reduce energy consumption
  - Total Energy = $E_{dynamic}$ + E_{static}
  - Dynamic Energy directly proportional to V_{DD}^2
  - E_{static} proportional to (V_{DD}* Leakage current * T_{clk})
- Optimum reduction of V_{DD}
  - Near-Threshold region
  - Lower limit of V_{DD} in commercial applications is ~70% of nominal [1]
  - Only 2X Energy Difference from Sub-Threshold
  - 10X Delay Difference from Sub-Threshold
  - Still >6X Energy Reduction as compared to Nominal

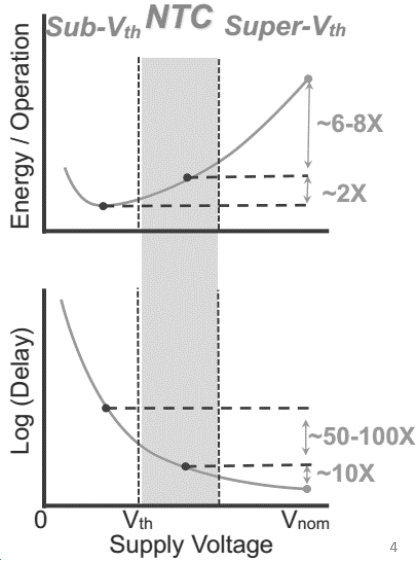

[Fick 2012] http://web.eecs.umich.edu/~mfojtik/fick_isscc2012_slides.pdf

4

Limitations of NTC: Soft Errors

- Soft Errors in logic datapath
  - Cause: Radiation-induced transient charge within a logic path which is ultimately latched by a F/F
  - More-than-ECC (Error Correcting Codes) needs to be done to mitigate soft errors for logic
- Soft Error Rate (SER) for logic at *NTV* is shown experimentally to be comparable to the SER for memory circuits [2]
  - Critical charge Q_{crit} needed to cause a failure decreases as V_{DD} is scaled. The SER has an exponential dependence on critical charge.
  - For 40nm and 28nm nodes, SER doubles when V_{DD} is decreased from 0.7V to 0.5V
- Soft Error masking mechanisms for logic paths
  - Logical Masking: fewer gate in critical path to regain lost throughput, less chance of the pulse being masked by logical computation of other gates in the path.
  - Electrical Masking: large pulse transients are created, as compared to supply voltage
  - Latching-window masking: lowered operating frequency has positive impacts here
- Non-planar devices offer a means to reduce SER.
  - 22nm Tri-gate technology is shown to reduce neutron and alpha particle induced SER at nominal voltage by 4-fold and 10-fold respectively compared to a 32nm planar process [3]
- Reduced pipeline depths, technology scaling, and NTV can be anticipated to have detrimental effects on logic SER

5

Adoption of NTC for Embedded applications

- A new direction for highly-reliable energy-efficient Embedded Processors/Chips
  - Energy-Efficiency → *NTC*
  - High-Reliability → *Spatial/Temporal Redundancy*
- Spatial Redundancy used in *mission-critical* applications for *resilience* as spare components help to tolerate failures [4],[5]
  - Harsh environments: autonomous vehicles, satellites, etc.
  - It is possible to reduce overhead by protecting only the critical components
- Temporal Redundancy (Repeated Execution) is also effective for soft error masking, however, performance loss is massive.
  - Suitable for area-constrained applications
- **Spatial:** *N*-Modular Redundancy (*N*-MR) and majority voting
  - System operates correctly as long as majority of the modules are functioning
  - Typically, *N=3* is employed: referred to as TMR systems

6

Soft Error Masking at NTV

- SER in logic paths can be reduced by schemes such as gate-sizing [6] and dual-domain supply voltage assignment [7]
  - Harden components which are more susceptible to soft errors. For instance, logic gates near the flip-flop
  - Difficult to provide comprehensive coverage. For instance, dual-domain voltage assignment is only able to reduce SER by 33.45%
- TMR provides comprehensive masking against soft-errors
  - Most soft-errors are mask-able or diagnosable
  - The probability of a non-diagnosable error is very low, i.e. what is the probability of majority instances producing identical and invalid outputs?
  - Temporal or Spatial Multiple Bit-Upset (MBU) should generate (with high probability) a diagnosable error

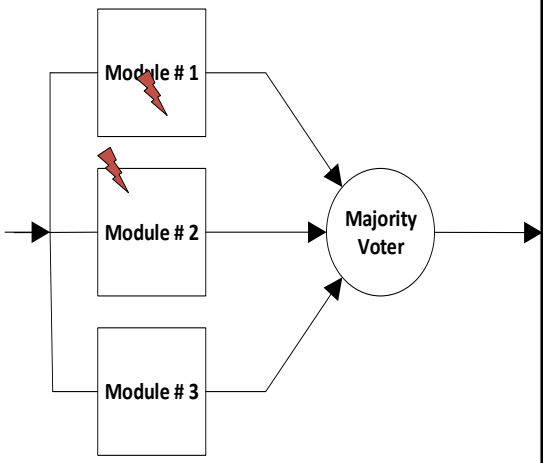

7

Related Work: Impact on Commercial Systems

- Variable Strength ECCs have been employed for reliable cache operation under aggressive voltage scaling [8]
- For processor caches operating at NTV, TMR is employed as a low-complexity means for improved resilience as compared to ECC schemes [9]
- Employing Modular Redundancy for High-Performance Computing (HPC) systems can significantly increase compute node *availability* [C. Engelmann et al. 2009]
  - HPC systems: decreasing MTTF, increasing MTTR due to scaling
  - Checkpoint and Restart is too costly (for complex HPC applications, increasing volume of state information needs to be saved)
  - Employing compute-node (processor(s), memory module(s), network interface) level redundancy permits to tradeoff individual component reliability by a factor of 100-100,000 → *$ Less Expensive $*

8

Limitations of NTC: Process Variations

- Near-Threshold Computing provides *Energy-Efficiency*
  - >10X Performance Loss → Parallelization [11], Device optimization [1]
  - (add-on) 5X Impact of Performance Variation → *Cost of Design Margins?*
- Nanoscale CMOS devices have Performance variability caused due to manufacturing-induced Process Variations (PV) [12].
  - For example, Random Dopant Fluctuations (RDF) are due to implanted impurity fluctuation and cause local variation (intra-die) in the threshold voltage of the transistors → *Increase in Delay Margins*
  - Impact of Technology Scaling: RDF magnified as number of dopant atoms is fewer so addition or deletion of just a few impurity atoms significantly alters transistor properties
- Operation near the threshold voltage of the transistors further exacerbates the process variability [1],[13]

Source: Borkar, Intel

Limitations of NTC: Delay Variations

- Delay measurements of FO4 Inverter Chains
- Implemented using PTM cards

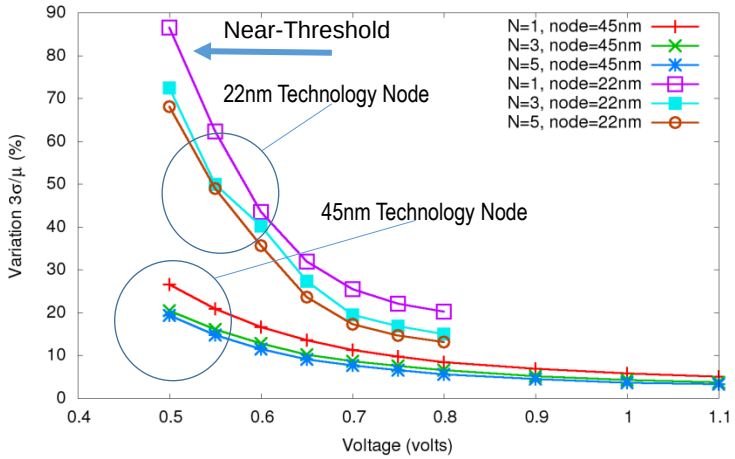

10

Modular Redundancy at NTV: What is the catch?

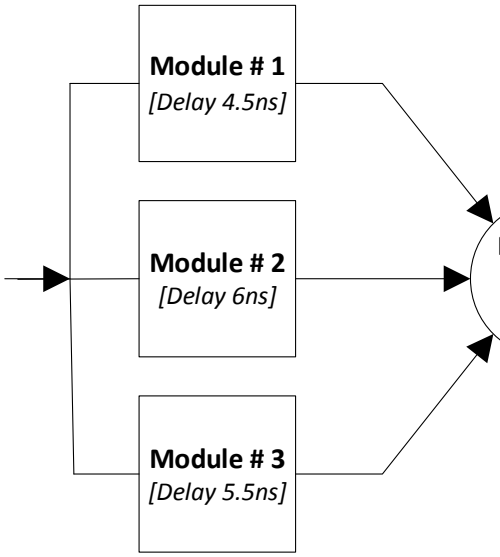

Need to consider the *worst delay* out of all *N* modules

$$\tau_{NMR} = \max_{1 \le i \le N} (\tau_i) + \delta$$

Clock = (1/6.5ns)

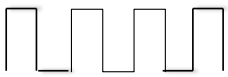

11

N-MR system Delay Distributions under PV

Delay Distributions (1000 arrangements each) at NTV of 0.55V with 45nm PTM model cards

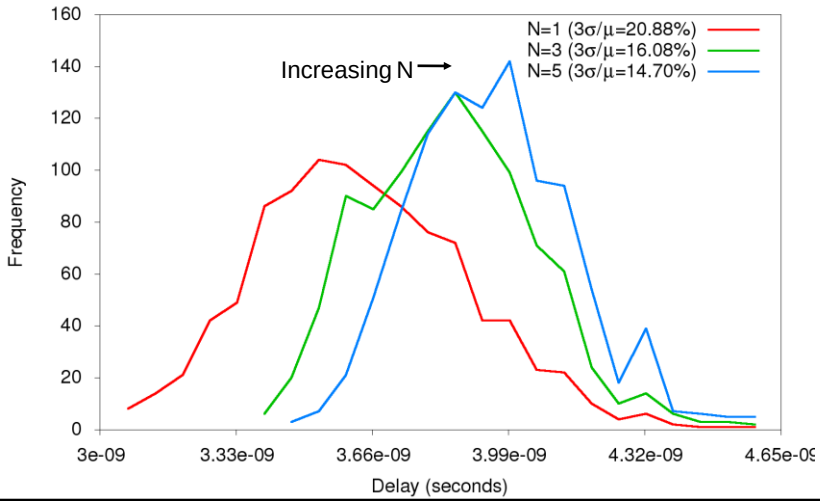

12

Performance of *N*-MR systems with scaled technology nodes

Mean delay difference of *N*-MR systems increases with voltage scaling down to Near-threshold region

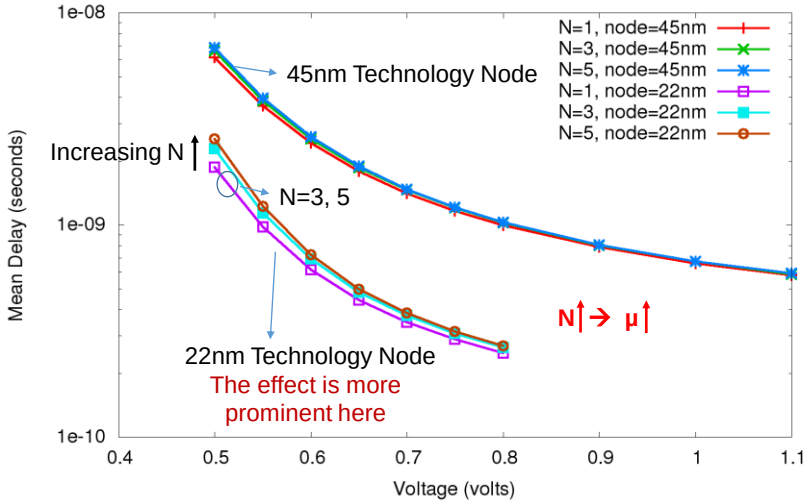

Performance of *N*-MR systems and variability

Delay Variations decrease with increasing N for *N*-MR systems

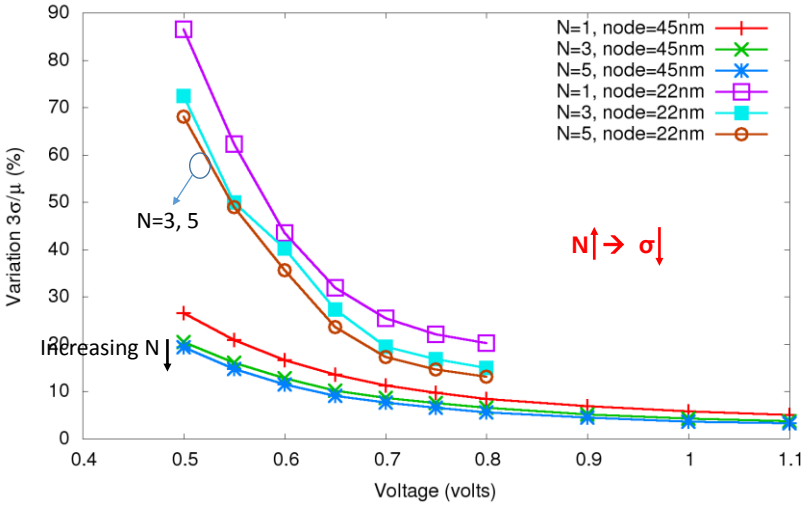

Reducing variability at NTV

- Variability is dependent on length of the critical path. More gates imply less variability [14]
- Type of logic gate utilized can impact variability

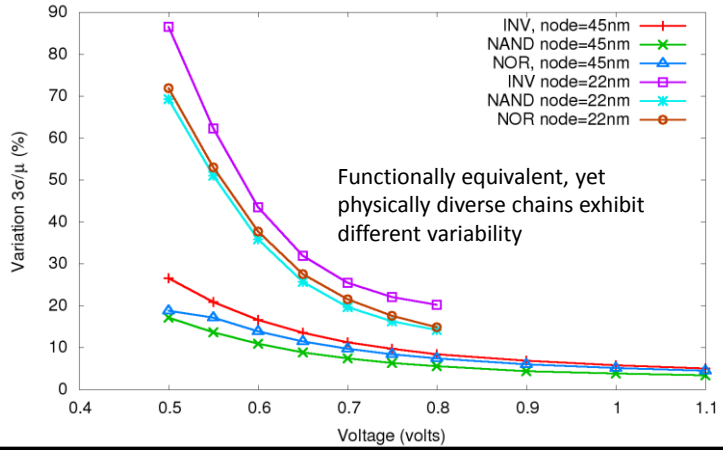

15

Reducing variability at NTV

- Develop a synthesis technique which realizes same function utilizing different gates with the goal of minimizing variability within given constraints

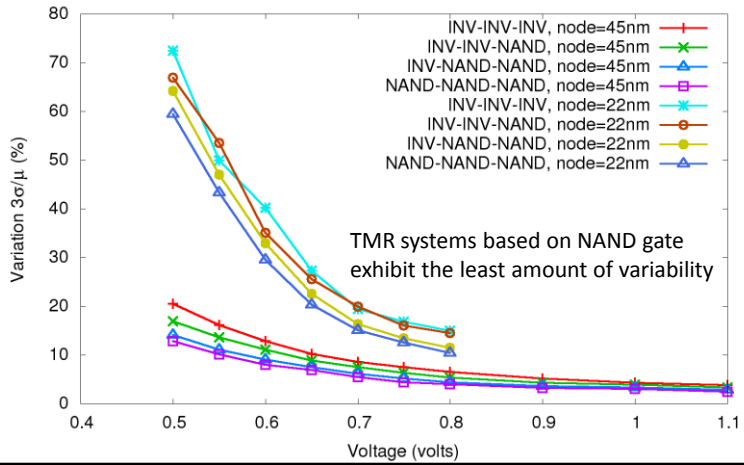

16

Future work: synthesizing variability immune circuit for NTV operation

- For our experiments with the inverter chains, the mean delays for NAND-based systems are higher than INV-based systems which outweighs any benefit of reduced variation.

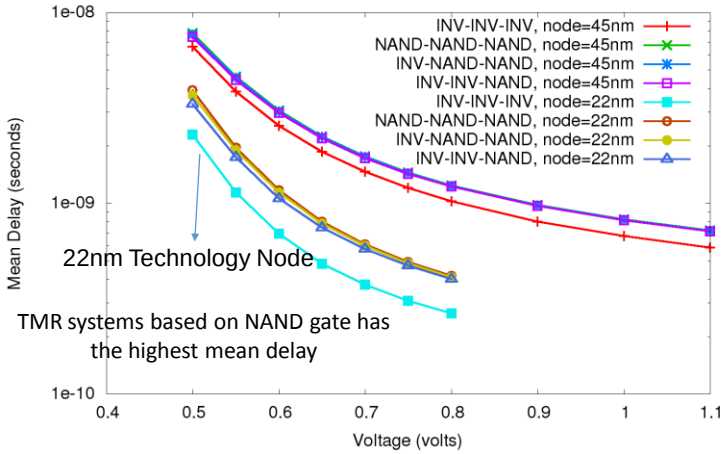

Energy Cost of Mitigating Variability

- "One-Time" Timing Guard-bands
  - Voltage and/or Frequency Margin [15]
- For a fixed V_{DD} of simplex system, how much voltage margin (ΔV_{DD}) needs to be added for *N*-MR system?
  - Left-shift the distribution of NMR system towards that of the simplex system
- Condition for same 99% Yield for *N*-MR system as compared to simplex system i.e., same delay [14]
  - (for $N \geq 3$): $\mu_{N\text{-}MR} + 3*\sigma_{N\text{-}MR} \leq \mu_{Simplex} + 3*\sigma_{Simplex}$
- How much energy overhead for *N*-MR system?
  - N-fold as mostly assumed with *N*-MR systems

Experimental Setup

- MCNC benchmark circuits `c880`, `i5`
- 45nm-based NanGate open source library [16]
- Synopsys Design Compiler used for synthesis
- Worst-case Test Vectors are generated using Synopsys TetraMax
- Synthesized netlists are imported into HSPICE for Monte-Carlo simulations
- Voter delay is not considered to make direct comparison to simplex systems

http://images.dailytech.com/nimage/22944_large_2007_08_monte_carlo.jpg

19

Experimental Setup

- At least 1000 Monte-Carlo iterations are performed
- μVth from Predictive Technology Model (PTM) cards
- RDF: σVth ranges from 25.9mV (45nm) to 59.9mV (22nm) [12]

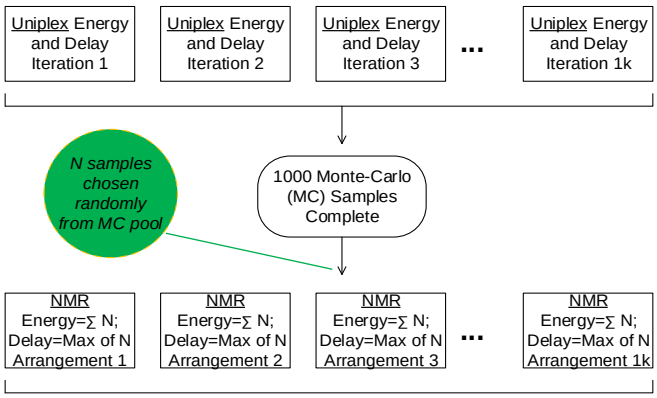

20

Results: Energy-Efficiency *vs* Reliability

- Energy measurements of FO4 Inverter Chains implemented using 45nm PTM cards

Operating Point	V_{DD} (V)	Normalized Energy	Normalized Delay
A	1.1	1X	1X
B	0.69	~1X	2.58X
C	0.545	~1.01X	7.15X

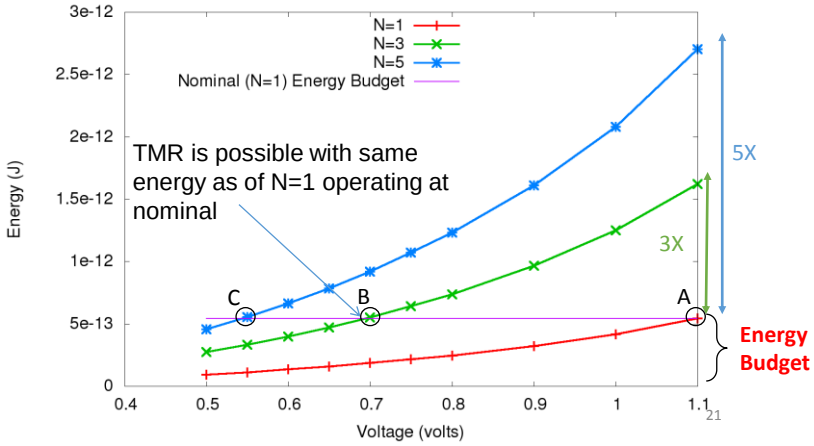

Energy Consumption of *N*-MR systems

- How much Voltage Margin (VM) expressed as ΔV_{DD} needs to be included for *N*-MR system?
- Even though *N*-MR systems exhibit higher mean delays, the reduced variance necessitates only a slight VM to meet the delay target of the simplex system
- On average 2mV increase in V_{DD} is satisfactory to operate a TMR arrangement based on `i5` circuit at comparable delay

Benchmark →	`c880`		`i5`	
Simplex, V_{DD} (NTV)	**N=3**	**N=5**	**N=3**	**N=5**
0.55 V	3.03X	5.06X	3.02X	5.05X
0.6 V	3.03X	5.05X	3.02X	5.04X
0.65 V	3.02X	5.04X	3.02X	5.04X
0.7 V	3.01X	5.03X	3.01X	5.02X

Results of VM: Increased Variability

- The mean delays for the 22nm (45nm) node with N = 3 & N = 5 are 1.16X (1.06X) and 1.24X (1.09X) the mean delay for a simplex system respectively at the same voltage of 0.55V
- The 22nm node-based 5MR system requires 3.94% more energy consumption than a similar configuration at 45nm
- At NTV of 0.5V for simplex arrangement, a 11mV increase in V_{DD} is required for the TMR arrangement for same performance

Technology Node →	**45nm**		**22nm**	
simplex, V_{DD} (NTV)	**N=3**	**N=5**	**N=3**	**N=5**
0.5 V	3.04X	5.07X	3.17X	5.30X
0.55 V	3.03X	5.05X	3.14X	5.27X
0.6 V	3.03X	5.04X	3.13X	5.26X
0.65 V	3.02X	5.03X	3.12X	5.23X
0.7 V	3.01X	5.02X	3.10X	5.16X

Conclusions and Future Work

- Redundancy provides a degree of freedom for increased reliability by diminishing the supply voltage
- Feasible when resulting increase in delay is tolerable
- Further study worthwhile to determine resilience provided by *N*-MR systems at NTV due to other noise sources such as:
  - Variation in supply voltage V_{DD} and Temperature
    - Expect a further variation of 2X [1] *(detrimental effect anticipated)*
  - Aging-induced variations [13]
    - Lower Voltage and Junction temperatures will lower aging effects such as Bias-Temperature Instability *(beneficial effect anticipated)*
    - Lower temperature and currents help to reduce interconnect defects due to Electromigration *(beneficial effect anticipated)*

24

References (1/2)

- [1] R. Dreslinski et al.,"Near-threshold computing: Reclaiming moore's law through energy efficient integrated circuits," Proceedings of the IEEE, vol. 98, no. 2, pp. 253–266, 2010
- [2] A. Dixit and A. Wood, "The impact of new technology on soft error rates," in Reliability Physics Symposium (IRPS), 2011 IEEE International, pp. 5B.4.1–5B.4.7, April 2011
- [3] S. Jahinuzzaman et al, "Alpha-particle Induced Soft Error Rates in 22nm Bulk Tri-Gate Technologies," 4th Annual IEEE Santa Clara Valley Soft Error Rate Workshop, 2012
- [4] J. Celis et al., "Methodology for designing highly reliable fault tolerance space systems based on COTS devices," in Systems Conference (SysCon), 2013 IEEE International
- [5] R. Al-Haddad et al., "Sustainable modular adaptive redundancy technique emphasizing partial reconfiguration for reduced power consumption," International Journal of Reconfigurable Computing, vol. 2011, 2011
- [6] Q. Zhou and K. Mohanram, "Gate sizing to radiation harden combinational logic," Computer-Aided Design of Integrated Circuits and Systems, IEEE Transactions on, vol. 25, no. 1, pp. 155–166, Jan 2006
- [7] K.-C. Wu and D. Marculescu, "Power-aware soft error hardening via selective voltage scaling," in Computer Design, 2008. IEEE International Conference on, Oct 2008,
- [8] C. Wilkerson, A. Alameldeen, and Z. Chishti, "Scaling the memory reliability wall," Intel Technology Journal, vol. 17, no. 1, pp. 18–34, 2013

25

References (2/2)

- [9] A. Seyedi, et. al, "Circuit design of a novel adaptable and reliable L1 data cache," 23rd ACM International Conference on Great Lakes Symposium on VLSI, pp. 333–334, 2013
- [10] C. Engelmann et al., "The case for modular redundancy in large-scale high performance computing systems," in Proceedings of the IASTED International Conference on Parallel and Distributed Computing and Networks (PDCN) 2009, vol. 641, Feb 2009.
- [11] E. Krimer et al., "Synctium: a nearthreshold stream processor for energy-constrained parallel applications," IEEE Computer Architecture Letters, vol. 9, no. 1, pp. 21–24, 2010.
- [12] Y. Ye et al., "Statistical modeling and simulation of threshold variation under random dopant fluctuations and line-edge roughness," Very Large Scale Integration (VLSI) Systems,IEEE Transactions on, vol. 19, no. 6, pp. 987–996, 2011.
- [13] H. Kaul et al., "Near-threshold voltage (NTV) design: Opportunities and challenges," in Proceedings of the 49th Annual Design Automation Conference, ser. DAC '12, 2012, pp. 1153–1158
- [14] S. Seo et al., "Process variation in near-threshold wide simd architectures," in Design Automation Conference (DAC), 2012 49th ACM/EDAC/IEEE, 2012.
- [15] U. Karpuzcu et al., "Coping with parametric variation at near-threshold voltages," Micro, IEEE, vol. 33, no. 4, pp. 6–14, 2013.
- [16] W. Zhao and Y. Cao, "New generation of predictive technology model for sub-45nm design exploration," in Proceedings of the 7th International Symposium on Quality Electronic Design, 2006.

26