TEXAS A&M UNIVERSITY

Use-it or Lose-it: Wearout and Lifetime in Future Chip-Multiprocessors

Hyungjun Kim,[1] Arseniy Vitkovsky,[2]

Paul V. Gratz,[1] Vassos Soteriou[2]

[1] Department of Electrical and Computer Engineering, Texas A&M University

[2] Department of Electrical Engineering, Computer Engineering and Informatics, Cyprus University of Technology

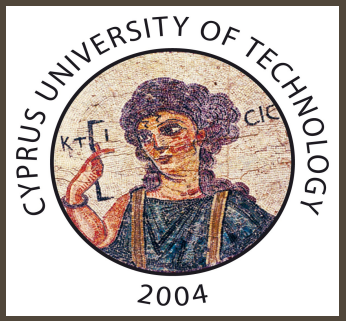

Introduction

- Moore's law scaling continues
  - Denard scaling ending
  - Power/performance trade-off driving core increases
  - Increasing core counts push complexity into interconnect
  - Interconnect becomes critical to CMP functionality

Intel Single-Chip Cloud Computer
48-core chip-multiprocessor (CMP)
2011

Moore's Law and Reliability

- ITRS: per-transistor failure rate must decrease 10X by 2023
  - "Manufacturable solutions *not* known"
- Failure mechanisms with time and use
  - *Hot Carrier Injection (HCI)*
  - *Negative Bias Temperature Instability (NBTI)*
- Transistor wear effects
  - Transistors slow w/ wear
  - Data dependent causes (roughly) inverse

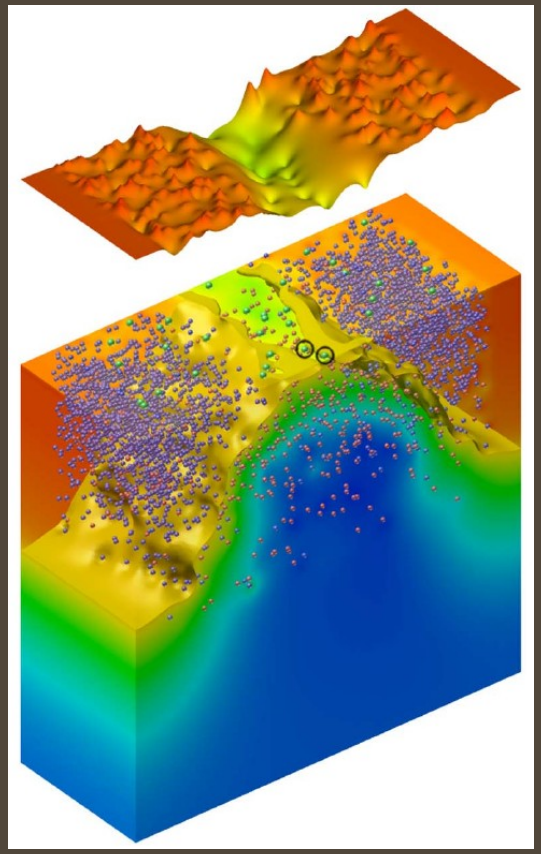

Charge trapping in pMOSFET after NBTI induced stress. [Brown et al., IEEE Trans. Electron Devices 2010]

Chip-multiprocessor Wearout

- Increasing core counts
- Each core accounts for less of system throughput

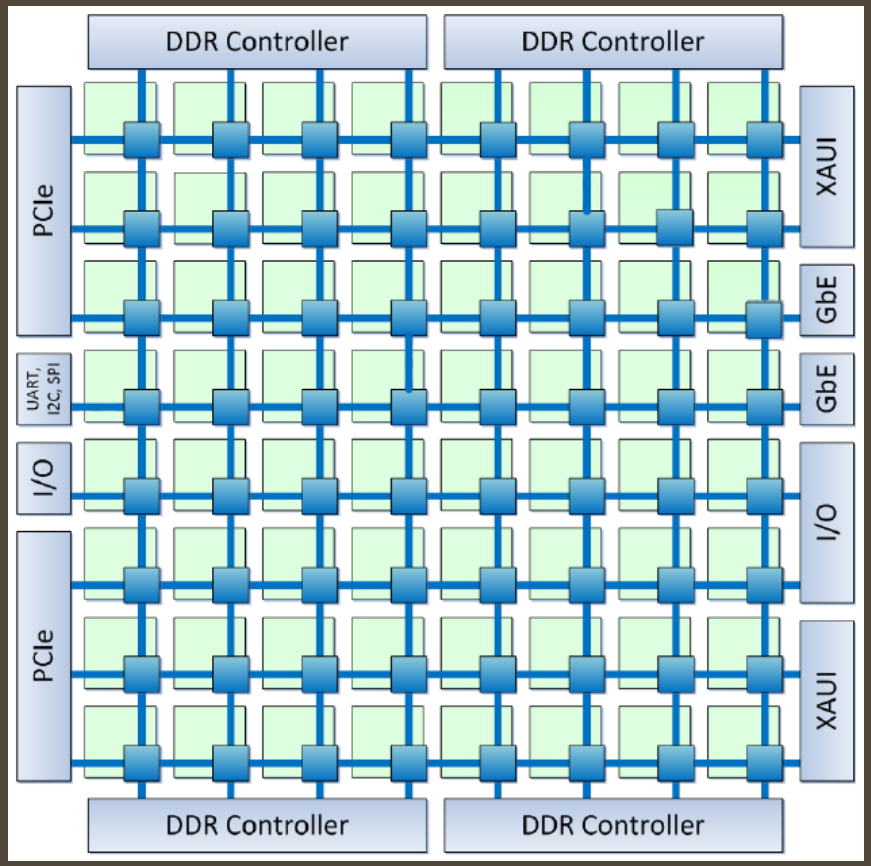

A 64-core Chip-Multiprocessor (CMP) with various peripherals interconnected via a 2-D Mesh

Chip-multiprocessor Wearout

- Increasing core counts
- Each core accounts for less of system throughput
- Core redundancy
  - Other cores can take over workload
  - Fault-tolerant task migration
  - Other micro-arch methods

- Problem solved?

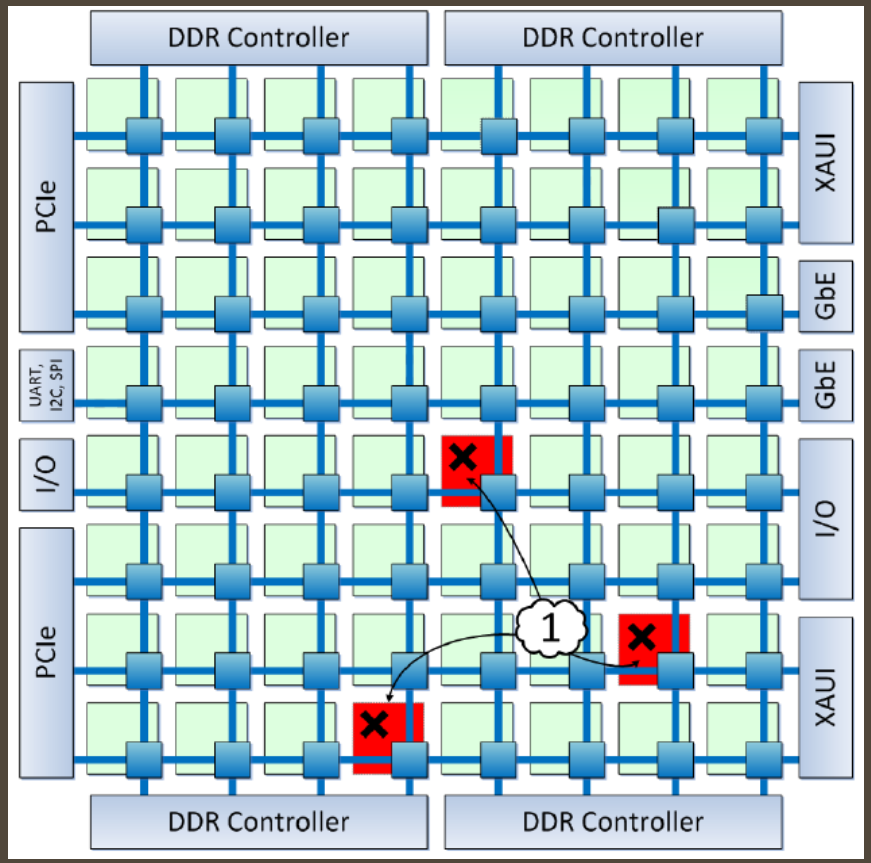

Scenario #1: Failure of three cores due to wearout, workload redistributed to other cores

Chip-multiprocessor Wearout

- Network-on-chip (NOC) complexity increases with scaling
- NoC critical to CMP operation

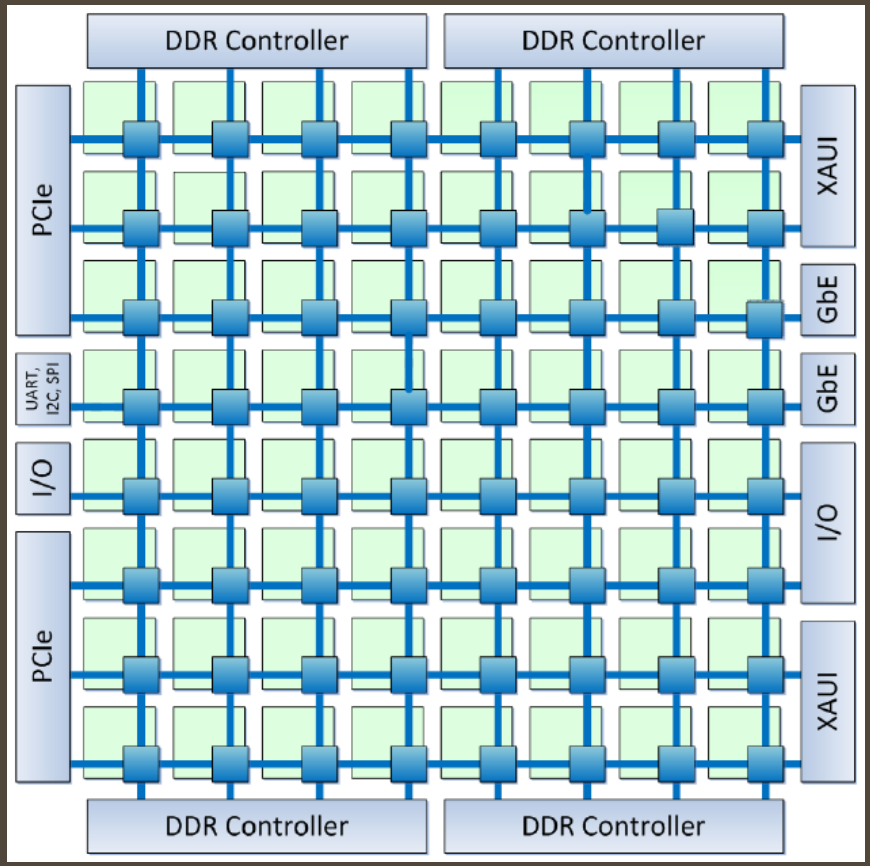

A 64-core Chip-Multiprocessor (CMP) with various peripherals interconnected via a 2-D Mesh

Chip-multiprocessor Wearout

- Network-on-chip (NOC) complexity increases with scaling
- NoC critical to CMP operation

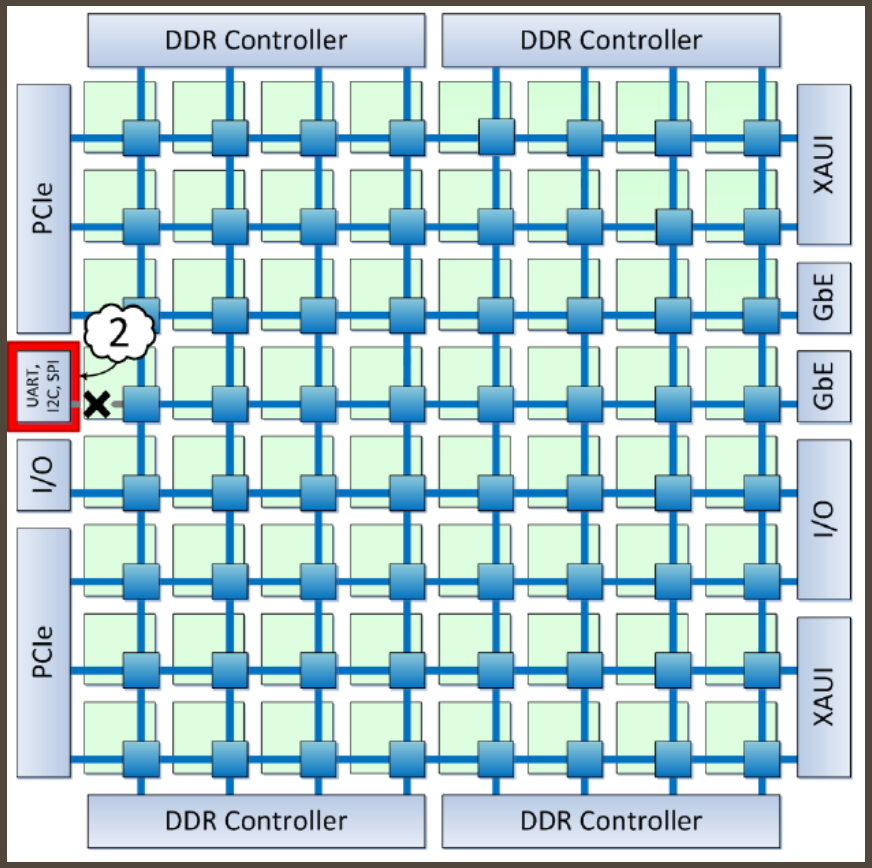

Scenario #2: A peripheral device disconnected from the CMP due to interconnect failure

Chip-multiprocessor Wearout

- Network-on-chip (NOC) complexity increases with scaling
- NoC critical to CMP operation

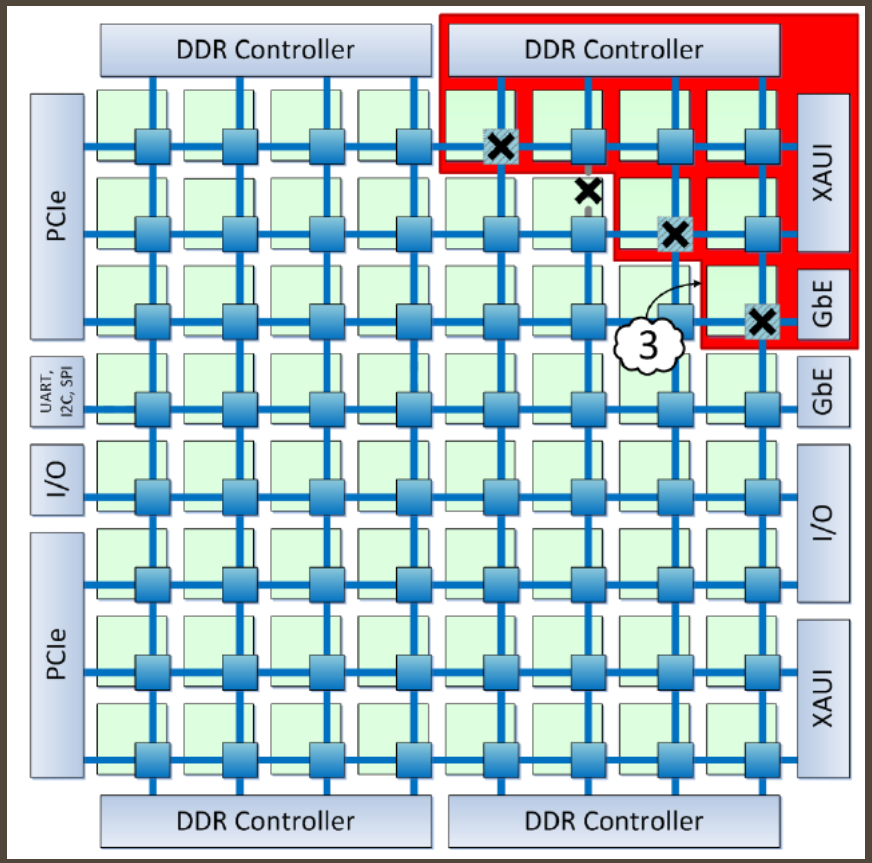

Scenario #3: Peripheral devices and entire NoC region disconnected from the network due to chained faults in links

Chip-multiprocessor Wearout

- Network-on-chip (NOC) complexity increases with scaling
- NoC critical to CMP operation

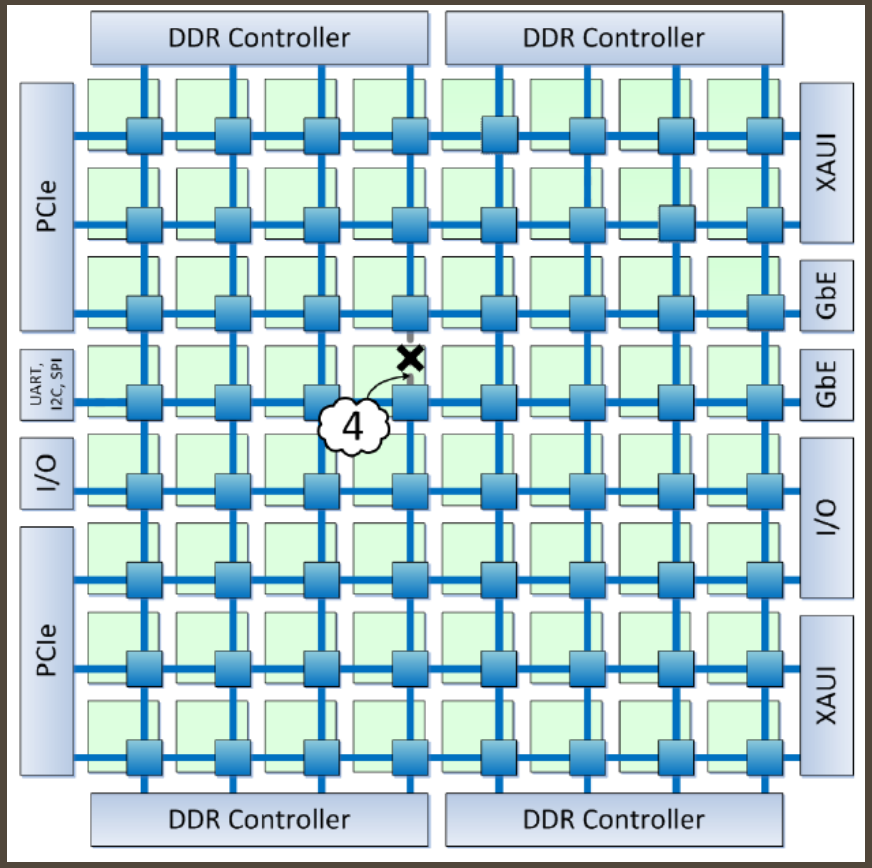

Scenario #4: A single link failure leading to routing deadlock

Chip-multiprocessor Wearout

- Network-on-chip (NOC) complexity increases with scaling
- NoC critical to CMP operation
- Single failure could incapacitate CMP!

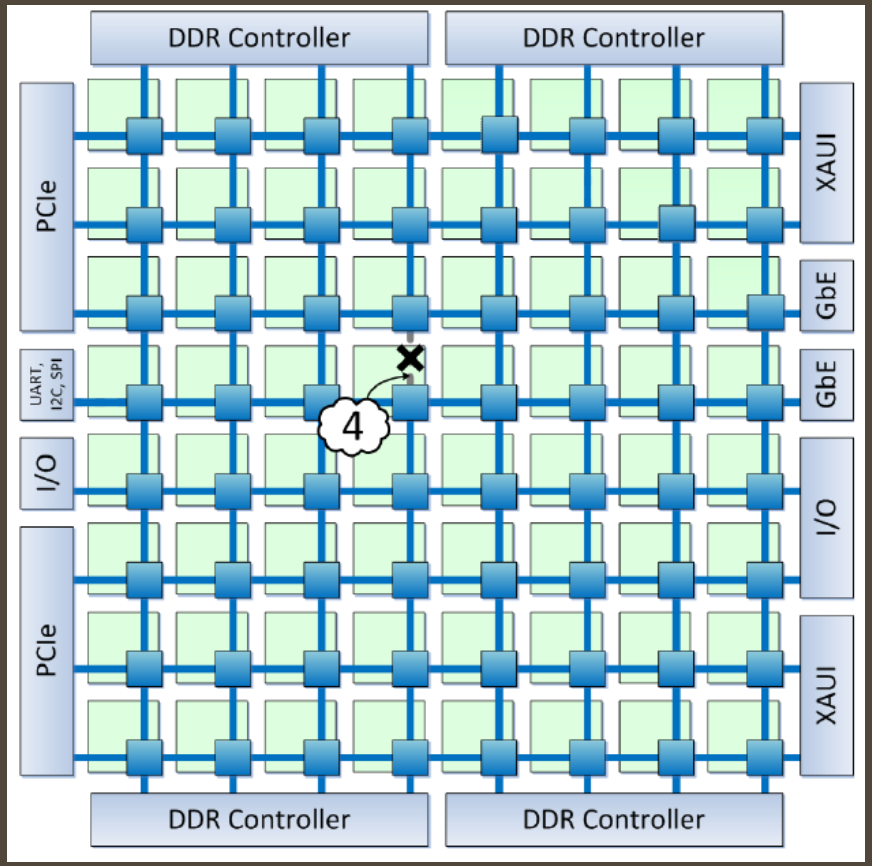

Scenario #4: A single link failure leading to routing deadlock

Chip-multiprocessor Wearout

Project goal:

- Develop techniques to mitigate wearout in the NoC
  - Analyze breakdown by mechanism
  - Characterize wear due to real workloads on routers and links
  - Develop wear-resistant router microarchitectures

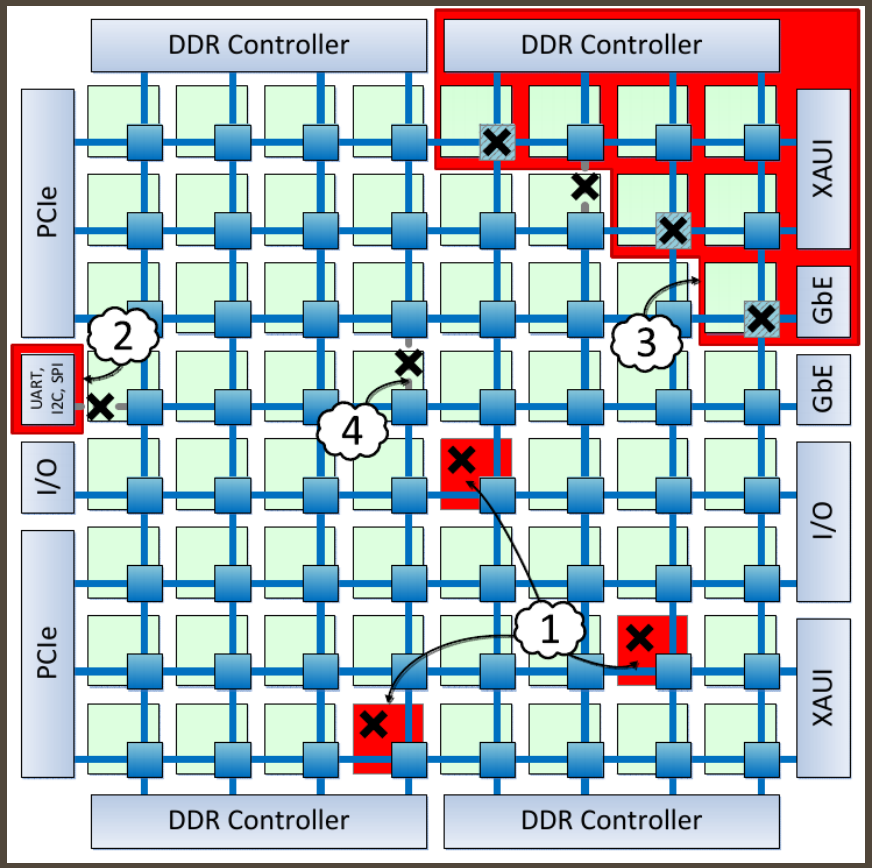

A 64-core Chip-Multiprocessor (CMP) with various peripherals interconnected via a 2-D Mesh, all failure scenarios illustrated

Outline

- Introduction
- Background
  - Failure mechanisms in modern CMOS
    - Hot Carrier Injection (HCI)
    - Negative Bias Temperature Instability (NBTI)
- Reliability characterization
- Wearout-resistant router microarchitecture
- Evaluation
- Conclusions

Failure mechanisms: Hot Carrier Injection (HCI)

- Cause:
  - When current flows through channel
  - Charge carrier gains sufficient energy to embed in gate oxide

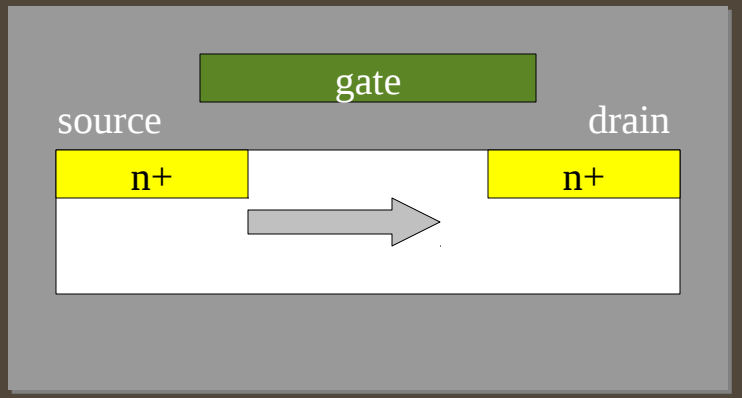

Failure mechanisms: Hot Carrier Injection (HCI)

- Cause:
  - When current flows through channel
  - Charge carrier gains sufficient energy to embed in gate oxide

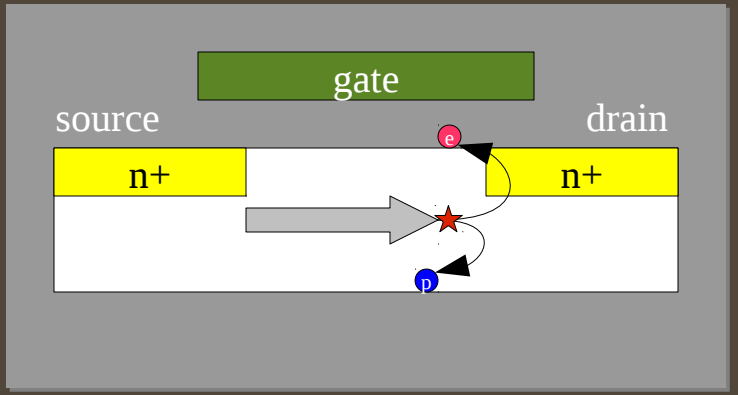

Failure mechanisms: Hot Carrier Injection (HCI)

- Cause:
  - When current flows through channel
  - Charge carrier gains sufficient energy to embed in gate oxide
- Effect:
  - Shifts Vt
  - “Slows” transistor switching speed

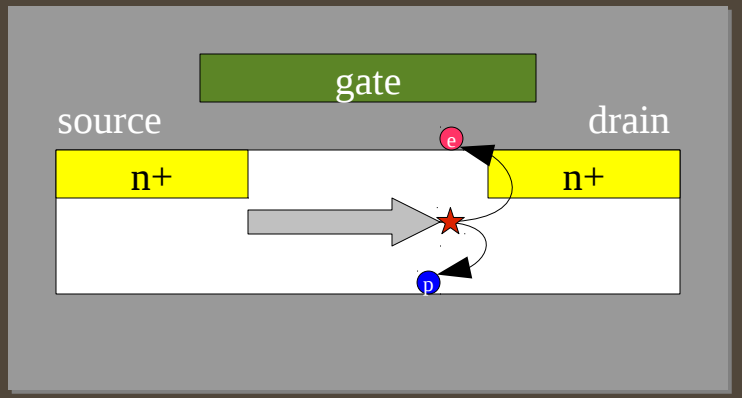

Failure mechanisms:

Negative Bias Temperature Instability (NBTI)

- Cause:
  - Hole induction breaks Si-H bonds at gate-oxide interface

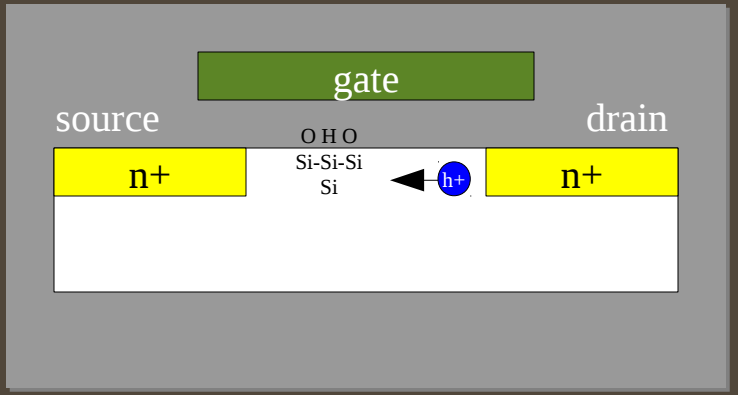

Failure mechanisms:

Negative Bias Temperature Instability (NBTI)

- Cause:
  - Hole induction breaks Si-H bonds at gate-oxide interface
  - H atoms form H^2 in oxide
  - H^2 diffuse to gate, trapping charge in gate oxide

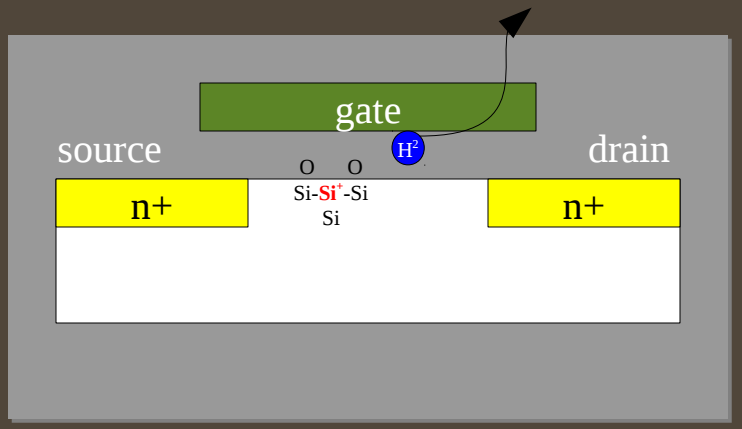

Failure mechanisms:

Negative Bias Temperature Instability (NBTI)

- Effect:
  - Shift in Vt
  - “Slows” transistor switching speed

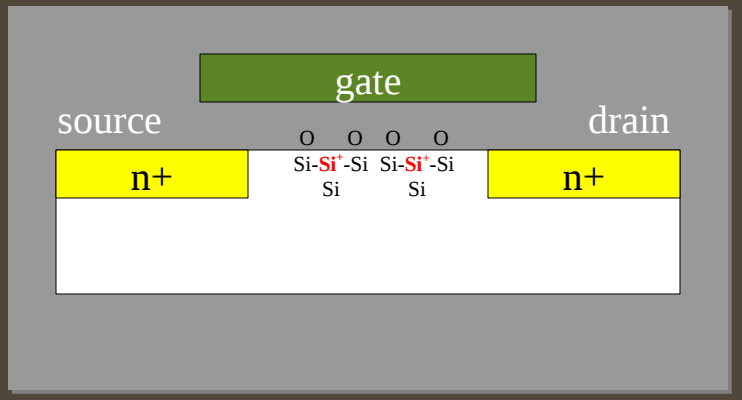

Failure mechanisms:
Negative Bias Temperature Instability (NBTI)

- Effect:
  - Shift in Vt
  - "Slows" transistor switching speed

- Effect worsens w/ high-K dielectric
- Much worse in MuGFETs

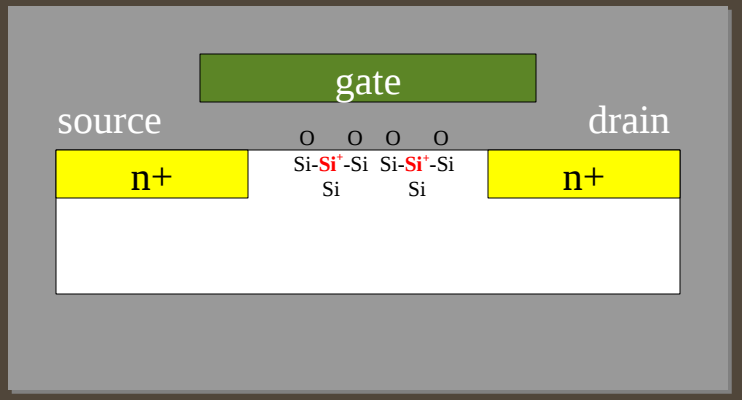

Failure mechanisms: Summary

TEXAS A&M UNIVERSITY

- Each mechanism occurs under different conditions
  - HCI proportional to current across channel
    - HCI increases with increased transistor activity factor
  - NBTI proportional to time gate held at 0V
    - NBTI increases with 0-skewed duty cycle
- Both mechanisms effect switching speed of transistor
  - Transistor does not “break” just slows down
  - Thus strongest impact on critical path
- Effect of usage on wear nearly opposite!
  - Increased activity → Increased HCI
  - Lower duty cycle → Increased NBTI

Outline

- Introduction
- Background
- Reliability characterization
- Wearout-resistant router microarchitecture
- Evaluation
- Conclusions

Router Microarchitecture

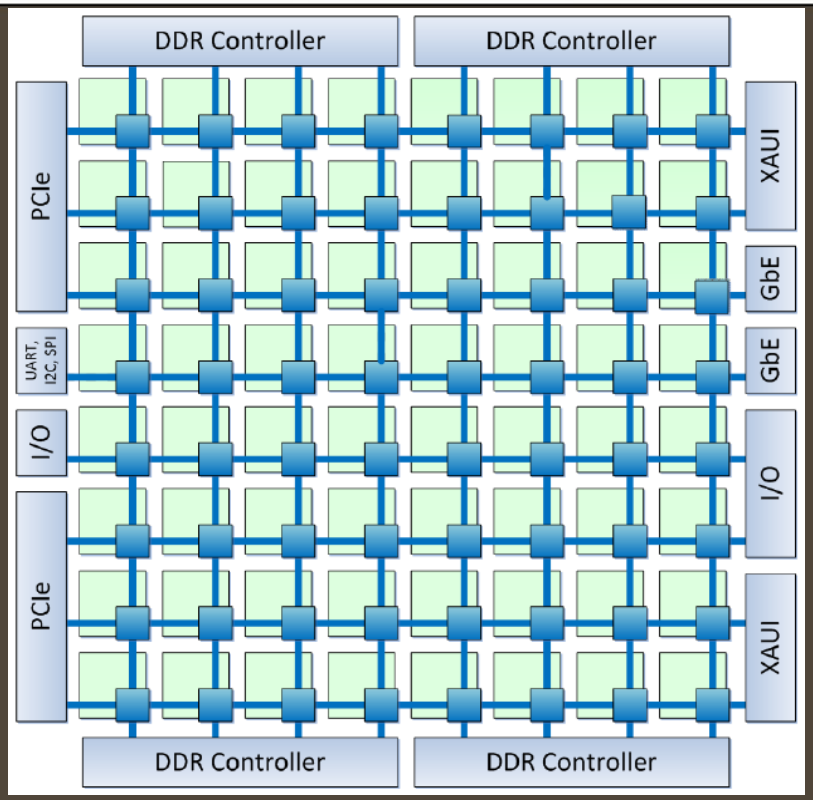

A 64-core Chip-Multiprocessor (CMP) with various peripherals interconnected via a 2-D Mesh

Router Microarchitecture

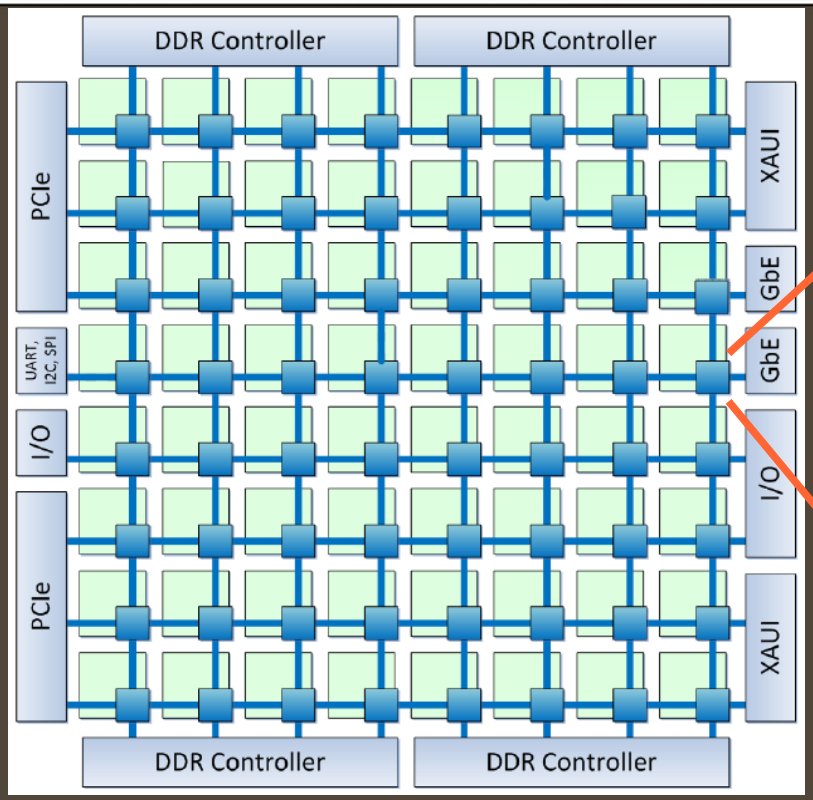

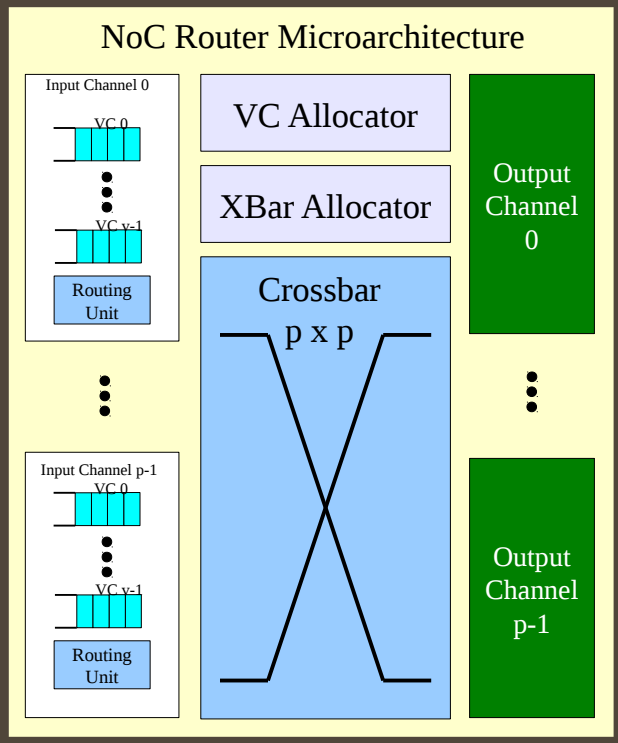

A 64-core Chip-Multiprocessor (CMP) with various peripherals interconnected via a 2-D Mesh

Router Critical Path

- HCI and NBTI slow transistors
- Wearout occurs first where timing critical
- Analyze critical path for susceptibility to wear

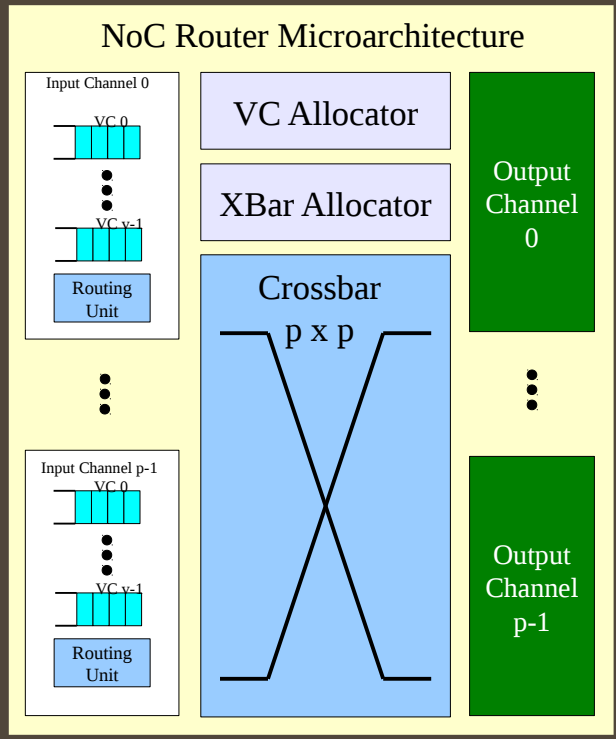

Router Critical Path

- NoC router critical path in 2nd pipeline stage
  - From input VC
  - Through VC and Xbar allocators
  - To 3rd stage pipeline latches

- Activity of wires along path correlated with injection rate of router
- Application dependent!

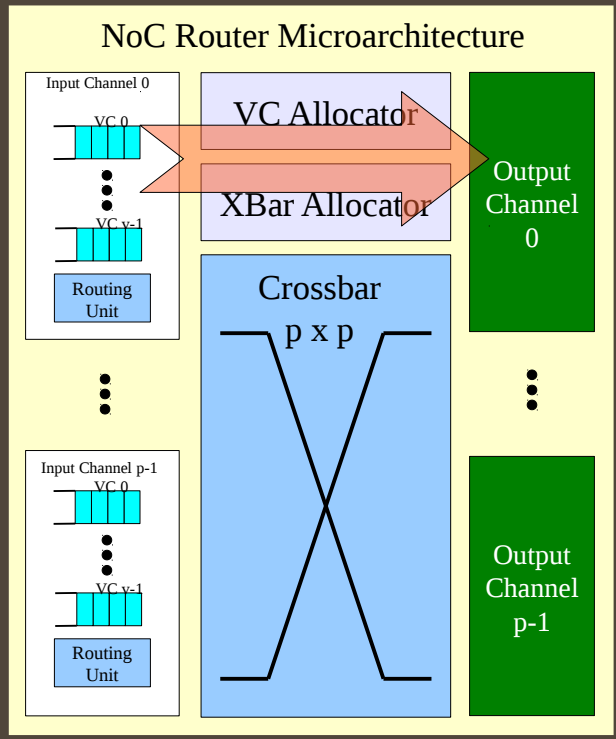

Router critical path typically lies through VC and Xbar allocation

CMP Workload Characterization

- Examined variance in per-router load
  - Incoming rate to each router individually
  - PARSEC workloads

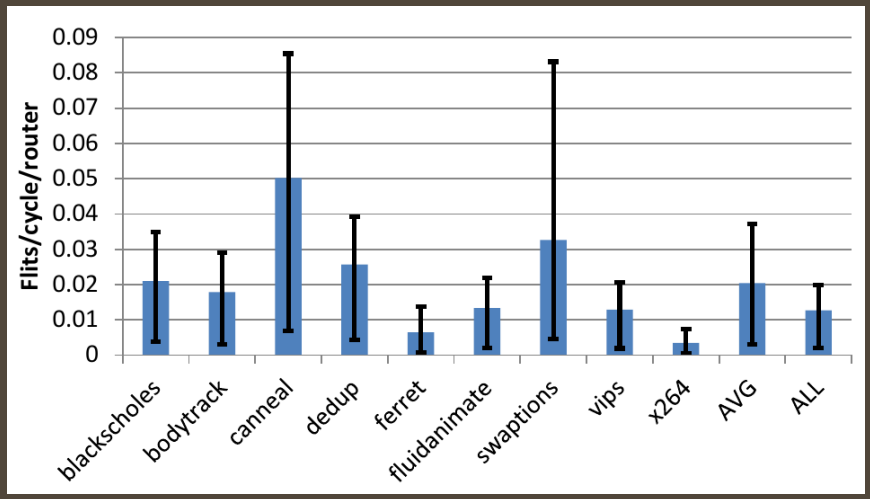

Load seen by each router under PARSEC suite benchmarks (incoming rate). Average (bar), minimum and maximum (whiskers).

CMP Workload Characterization

- Examined variance in per-router load
  - Incoming rate to each router individually
  - PARSEC workloads
- Great variance in load
  - Between applications
  - From router to router within network
- Implication: wearout is highly workload dependent

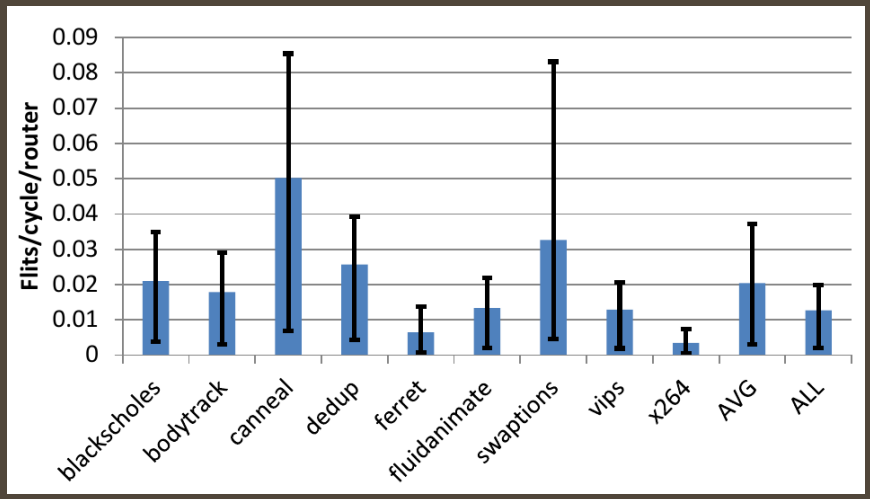

Load seen by each router under PARSEC suite benchmarks (incoming rate). Average (bar), minimum and maximum (whiskers).

Workload impact on wearout: HCI – Activity on critical path

Examined activity factor of router critical path to determine HCI impact

- Some sensitivity to injection rate
- Even for extreme injection rates, activity factor remains low

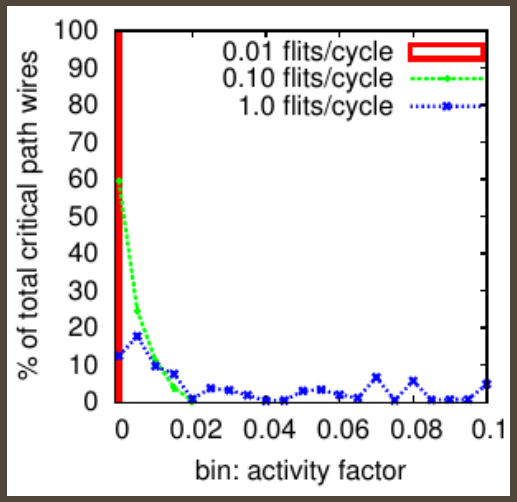

Activity factor histogram for critical path nodes within router, under random injection of varying rate. Critical path nodes defined as nodes on all paths with latency within 10% of clock period.

Workload impact on wearout: HCI – Activity on critical path

Examined activity factor of router critical path to determine HCI impact

- Some sensitivity to injection rate
- Even for extreme injection rates, activity factor remains low

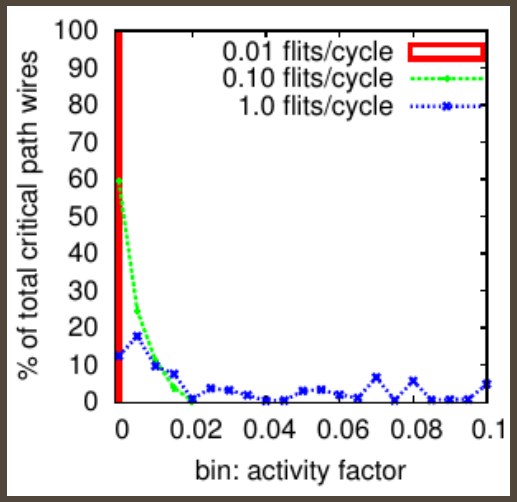

Activity factor histogram for critical path nodes within router, under random injection of varying rate. Critical path nodes defined as nodes on all paths with latency within 10% of clock period.

Implication: Activity factor on critical path low, and insensitive to workload – HCI impact low

Workload impact on wearout: NBTI – Duty cycle on critical path

Examined duty cycle of router critical path to determine NBTI impact

- Duty cycle poor for a high percentage of critical path nodes
- Highly sensitive to injection rate
  - Duty cycle improves as injection rate increases

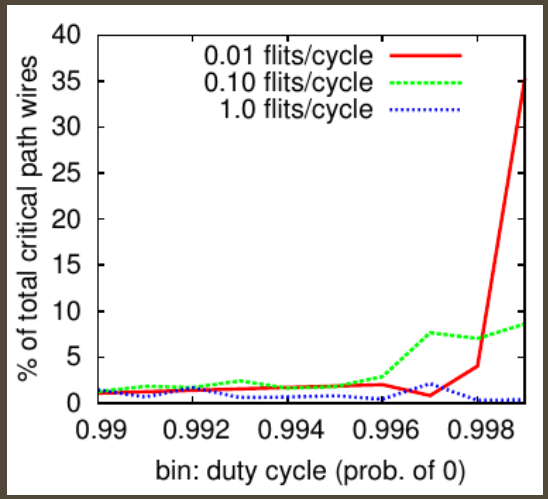

Duty cycle (fraction of time at 0) histogram for critical path nodes within router, under random injection of varying rate.

Workload impact on wearout: NBTI – Duty cycle on critical path

Examined duty cycle of router critical path to determine NBTI impact

- Duty cycle poor for a high percentage of critical path nodes
- Highly sensitive to injection rate
  - Duty cycle improves as injection rate increases
- Implications:
  - NBTI induced wear is likely
  - NBTI stress inversely proportional to load!

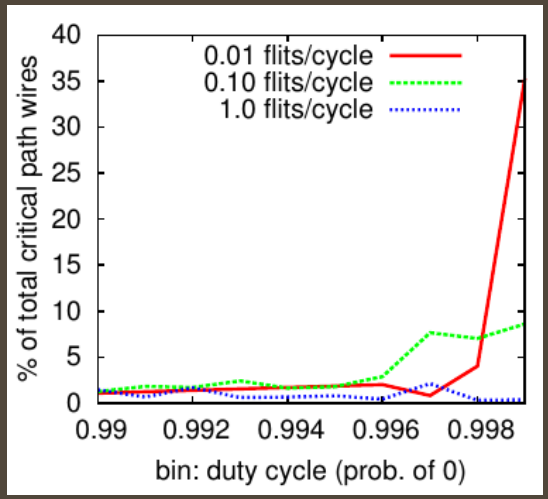

Duty cycle (fraction of time at 0) histogram for critical path nodes within router, under random injection of varying rate.

Reliability characterization summary

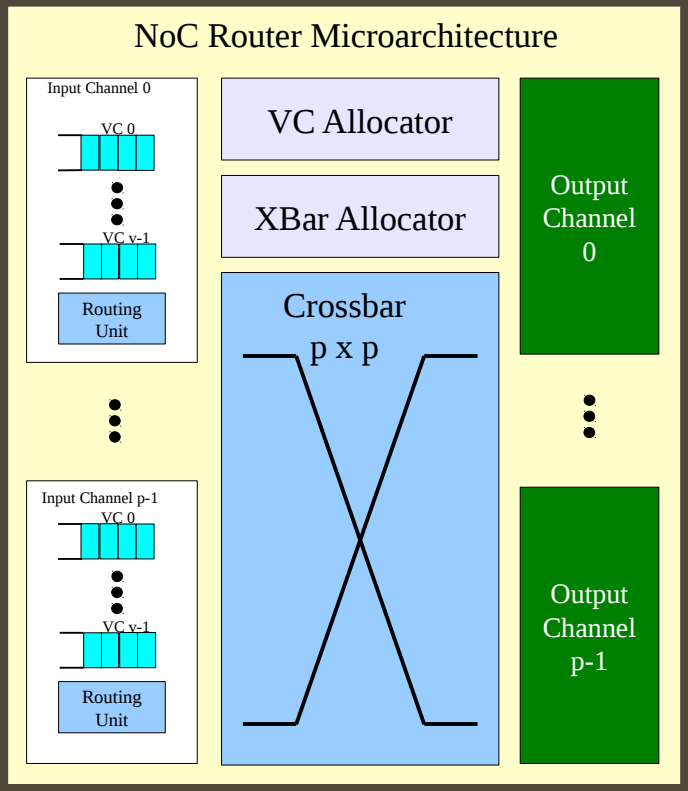

- Wear is injection rate dependent along the critical path
- Low injection rates cause poor duty cycle and accelerated NBTI
- Critical path primarily concerned with VC and Xbar allocation

Outline

- Introduction
- Background
- Related work
- Reliability characterization
- Use it or Lose it: wearout-resistant router microarchitecture
- Evaluation
- Conclusions

Addressing wear in the NoC router

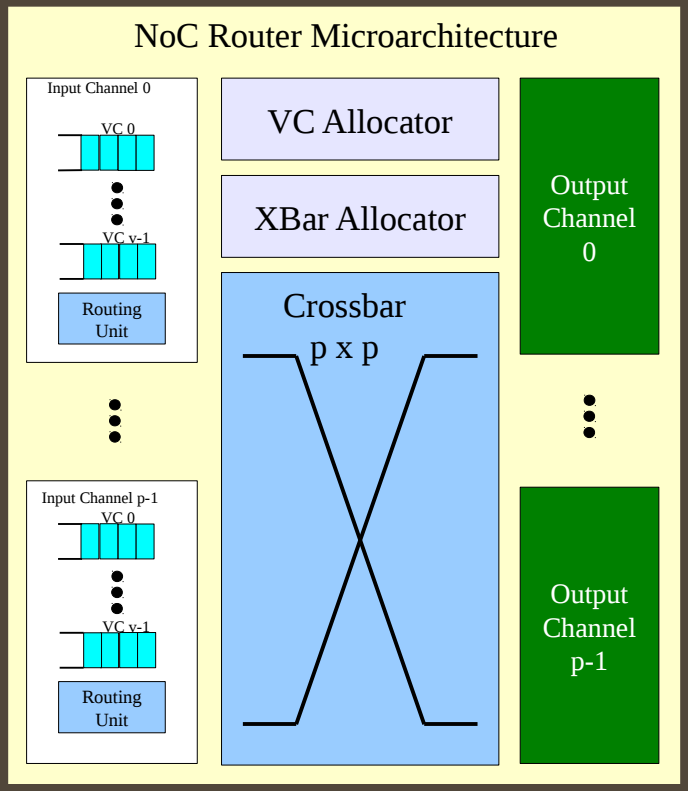

- Increasing injection rate artificially yields other problems
  - Power increases
  - HCI increases
  - Congestion
- Want to improve duty cycle w/o increasing activity factor

The *Use it or Lose it* router microarchitecture

TEXAS A&M UNIVERSITY

Use it or Lose it: Need mechanism to exercise critical path through the allocators or we will "lose" them

Goals:

1. Improve duty cycle by creating more time spent at "1"
2. Should not change state of router
3. Should not worsen critical path timing
4. Should not significantly increase activity factor
   - Power and HCI are activity factor dependent

The *Use it or Lose it* router microarchitecture

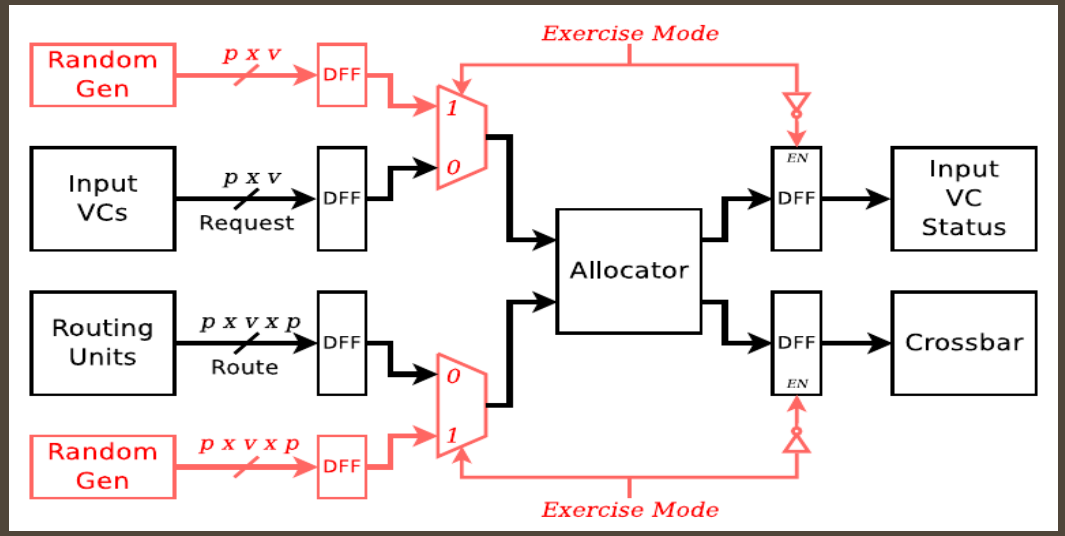

Create an “exercise mode” for critical path

1. Improve duty cycle by creating more time spent at “1”
   - Inject random data into allocators
   - Emulate higher injection rates
   - Hit corner cases

The *Use it or Lose it* router microarchitecture

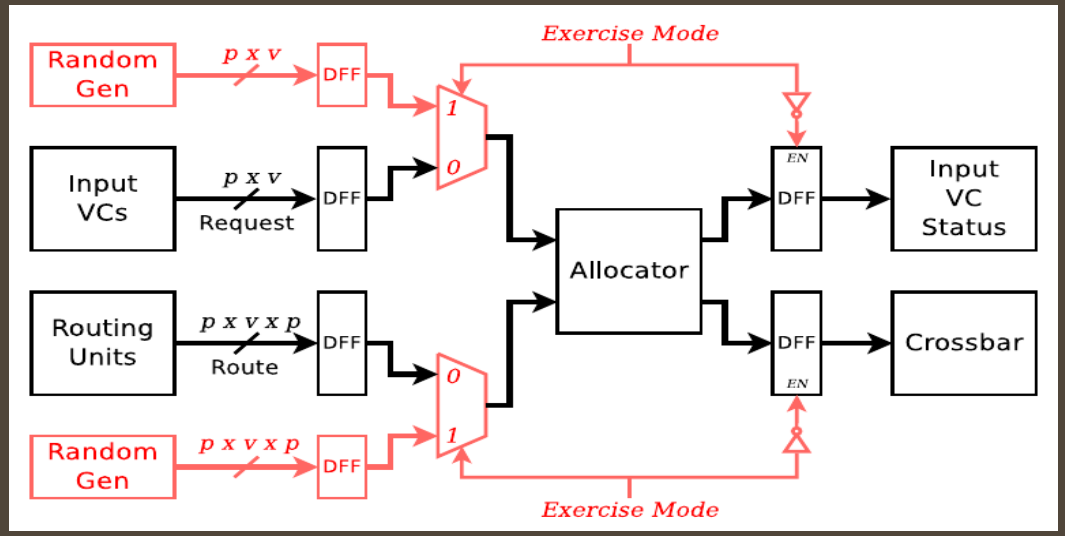

2. Should not change state of router
   - Only enabled when no packets ready in input FIFOs
   - Router state latches are disabled during exercise mode
3. Should not worsen critical path timing
   - Only one mux added to actual path
   - All vector generation and mode selection logic off critical path

The *Use it or Lose it* router microarchitecture

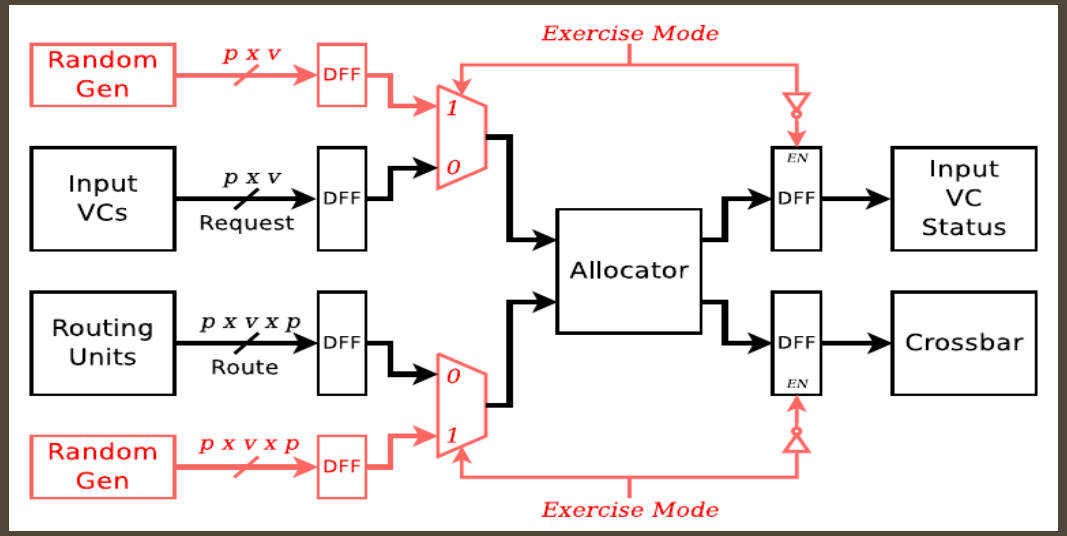

4. Should not significantly increase activity factor
   - Only paths with poor duty cycle on critical path are modified
   - Vector generator changes vector slowly, once every 128 cycles
   - Minimal impact on router activity factor

Random Vector Generation

- ROM containing pre-generated random vectors
  - Low overhead
  - Examined range of ROM sizes
    - ROM sizes >16 produced little return
  - Examined range of ROM vector rotation periods
- Examined other methods to generate random vector (see paper)

Outline

TEXAS A&M UNIVERSITY

- Introduction
- Background
- Related work
- Reliability characterization
- Wearout-resistant router microarchitecture
- Evaluation
- Conclusions

Implementation

- Baseline and Wear-resistant router synthesized in 45nm TSMC process tech
  - Results scaled to 22nm
- Low overheads of technique
  - ~1% increase in router power consumption
  - <1% increase in router area

Methodology

- Post-synth router models developed for Baseline and Wear-resistant
- Post-synth router models simulated under PARSEC-like workloads
  - 64-node, 2-D Mesh (8x8)
  - Five-port, three-stage, 2-D mesh router
  - Four-VCs per port, each four flits deep
- Activity factor and duty cycle of wires along router critical paths extracted

Results: Lifetime Increase Under Synthetic Workload

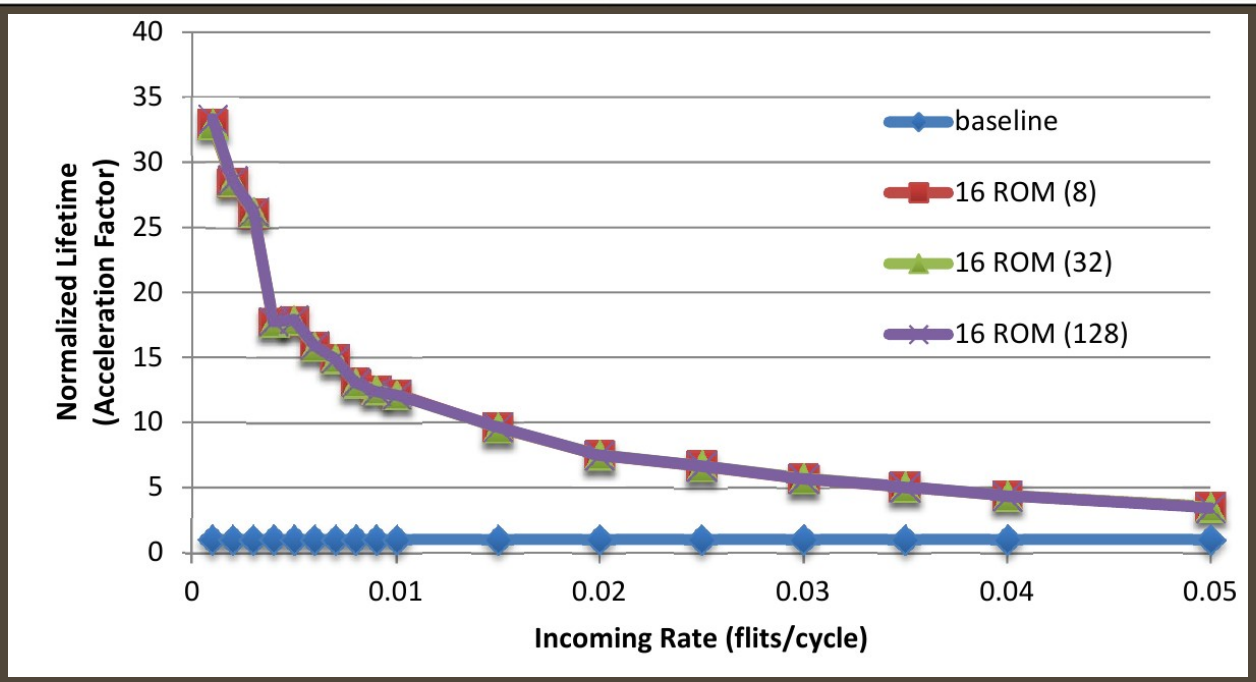

Notation: "16 ROM (8)": 16 entry ROM, rotated every 8 cycles.

Lifetime of *Use it or Lose it* router experiencing NBTI wear under various injection rates normalized against the baseline router.

- Dramatic improvement in lifetime
  - Improvement best at low loads when wear is worst
- Insensitive to rotation period

Results: Lifetime Increase Under PARSEC workloads

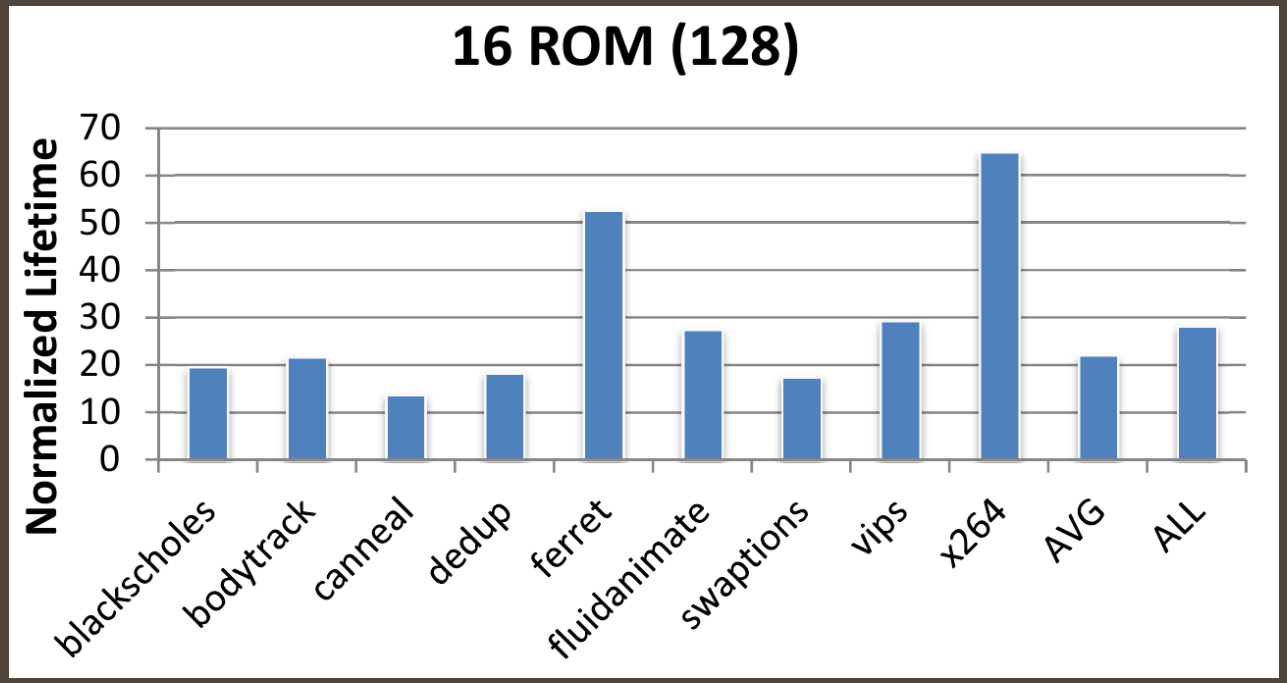

Lifetime improvement of *Use it or Lose it* router under PARSEC workloads normalized against baseline router.

- 10-65X increase in lifetime over baseline
  - Geo-mean of 22X improvement across benchmarks
  - Applications with routers that experience lowest minimum load see greatest gain

Results: Impact on Activity Factor

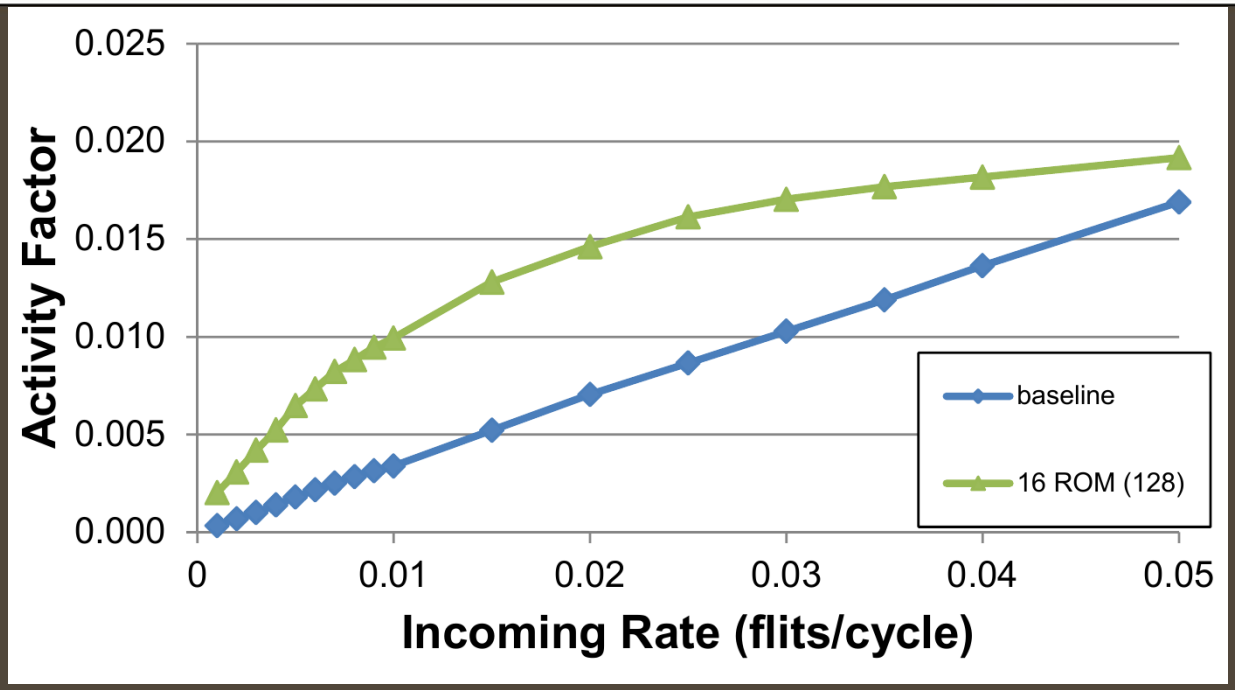

Activity factor for baseline and *Use it or Lose it* router under various incoming rate loads.

- Increase in activity factor only significant at low, < 0.05 flits/cycle, loads
  - Little impact on HCI

Results: Impact on Power

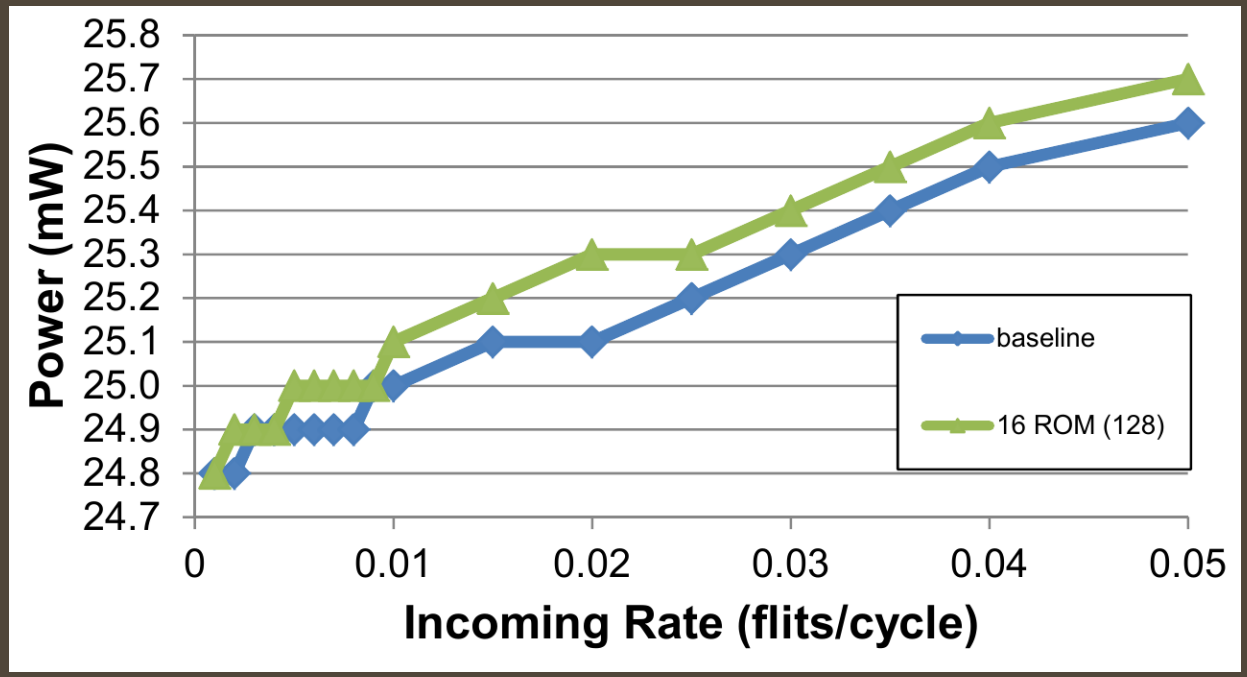

Power consumption of baseline and *Use it or Lose it* router under various incoming rate loads.

- Maximum of < 1% increase in router power consumption

Conclusions

- Wearout is a growing problem in future CMOS technologies
  - Over 10 years, 10x reliability improvement required
- Interconnect (NoC) critical point of failure in CMPs
- In the NoC, NBTI is dominant
  - Low activity of routers problematic
  - HCI less of an issue
- Introduced *Use it or Lose it* wear-resistant router microarchitecture
  - 22x improvement in lifetime
  - Insignificant overhead in energy or performance

TEXAS A&M UNIVERSITY

Questions?