High Performance Networking Hardware Design

Gokhan Memik
Center for Ultra-Scale Comp. and Information Security (CUCIS)
Electrical Engineering and Computer Science Department

NORTHWESTERN
UNIVERSITY

NORTHWESTERN UNIVERSITY

Problem and Overview

- Need for faster processing
  - Increasing link speeds
  - Changing Bottlenecks
    - Transmission >> processing => transmission ~= processing => transmission << processing
- Particularly important for distributed applications
- Traditional approach: higher levels of parallelism and specialization
  - Costly, does not scale, is not applicable
- Our approach: Vertical integration
  - Design systems (including applications, run-time systems, architectures, and circuits) in an integrated fashion

NORTHWESTERN UNIVERSITY

Vertical Integration

- Example: Correctness
- 100% Correctness
  - Consider all cases, which may never occur
- Correctness is overrated
- Networking
  - Higher levels (e.g., TCP) take care of errors
  - Processors can be more aggressive
- 2 to 4 times improvement with less than 0.1% error rate

Treat correctness as an objective, not a requirement

NORTHWESTERN UNIVERSITY

Overview

- Task 1: Design and analysis of fault vs. performance trade-offs (fault throttling)
  - E.g, overclocking
- Task 2: Measure the impact on a single router
  - E.g, architectural simulation
- Task 3: System or Network level implications
  - E.g, network simulation

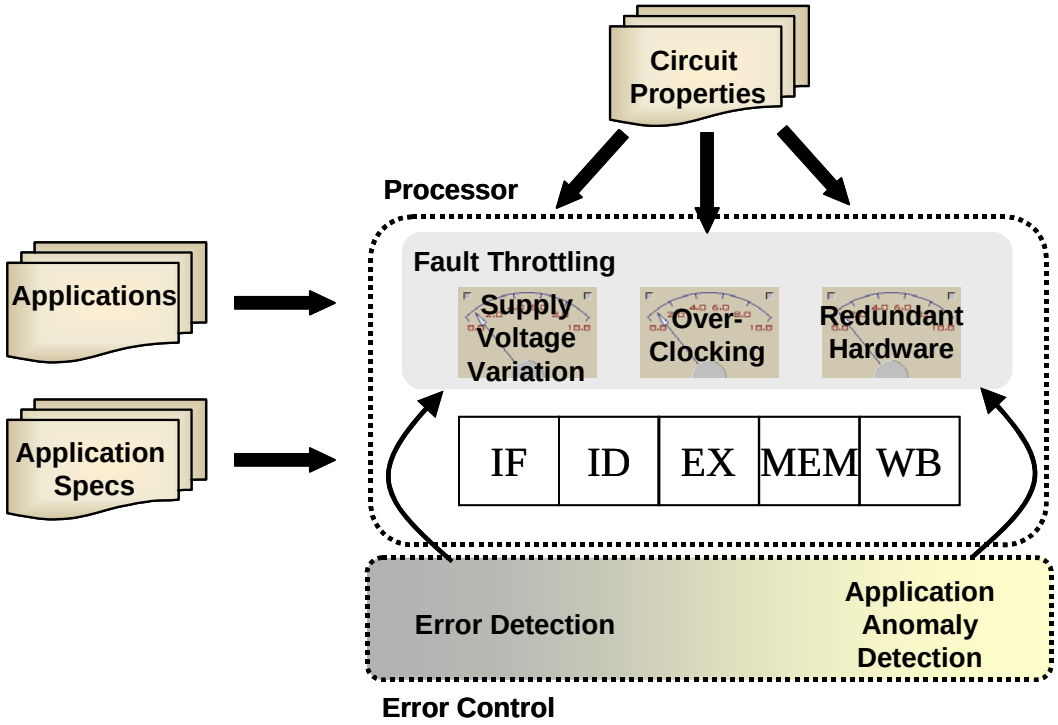

Overclocking (Task 1)

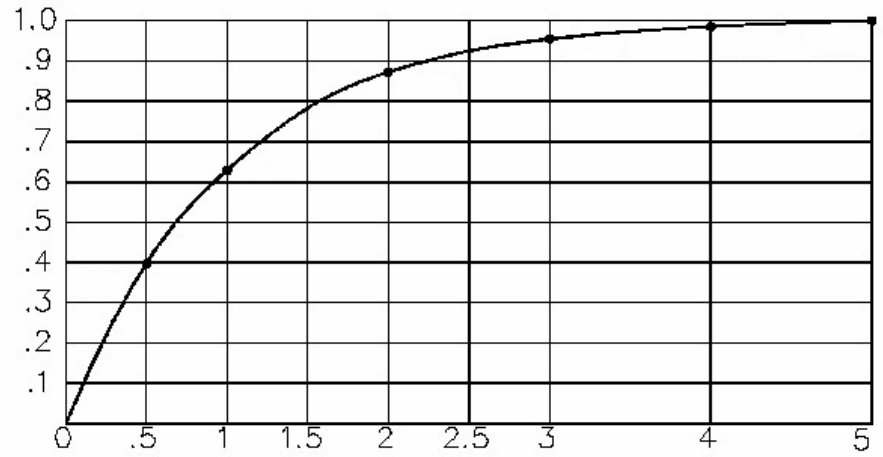

Voltage Swing vs. Time

- Voltage swing
  - Rapid increase at first
  - Slow increase later
- Overclocking

NORTHWESTERN UNIVERSITY

Not so fast!

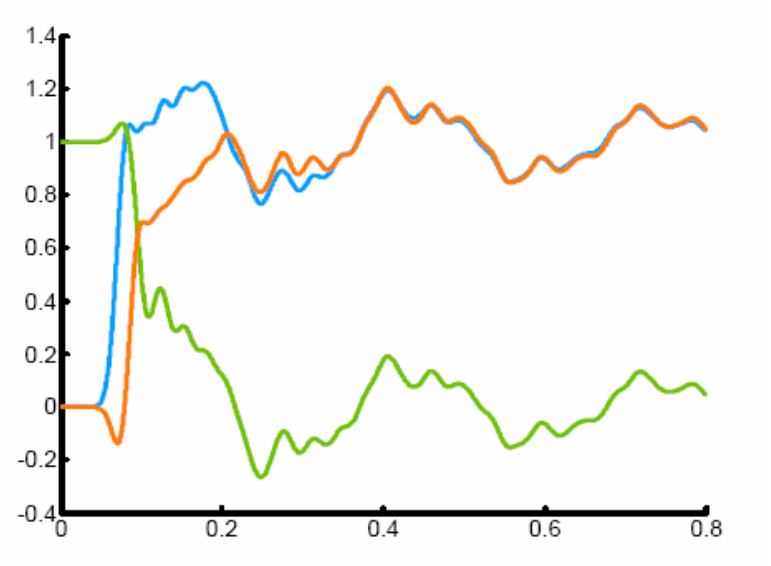

- Noise (inductive and/or capacitive)
  - Signal deviation
- Overclocking
  - Reduced immunity

Approach

- Analyze each component separately

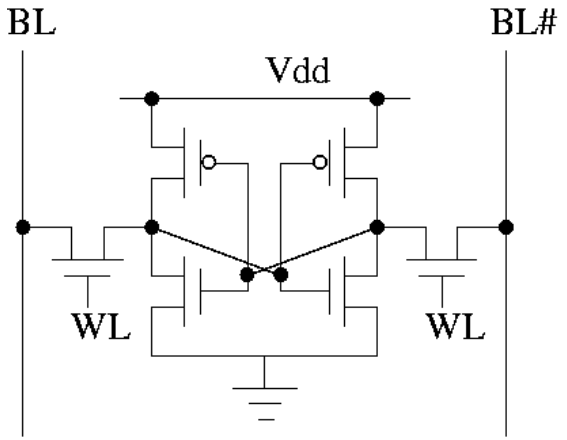

- 6-transistor SRAM cell
  - Input, clock, feedback loop

Modeling Faults

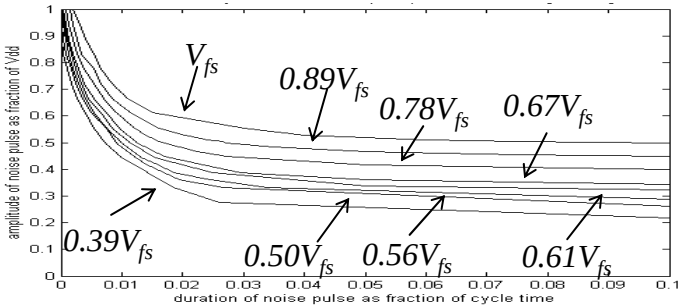

Noise immunity curves

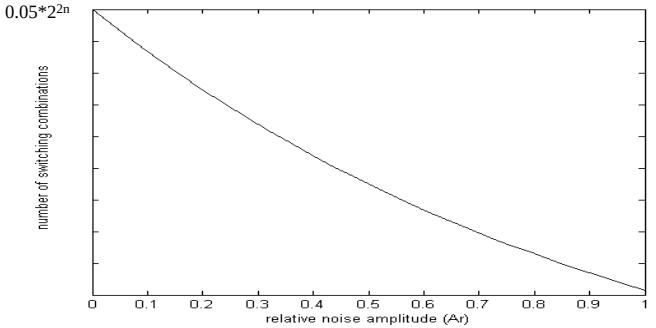

Noise amplitude for switching comb.

- Analyze the impact of noise on the feedback loop ⇨ Noise immunity curves
- Different noise amplitude probabilities
  - Check all switching combinations

$$P(A_r) = 28.8 * e^{-28.8 A_r}$$

Estimation model

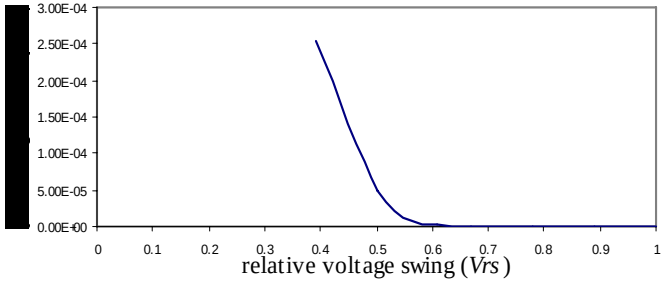

Fault probability versus voltage swing

- Fit distribution into immunity
- Combine it with voltage swing vs. time

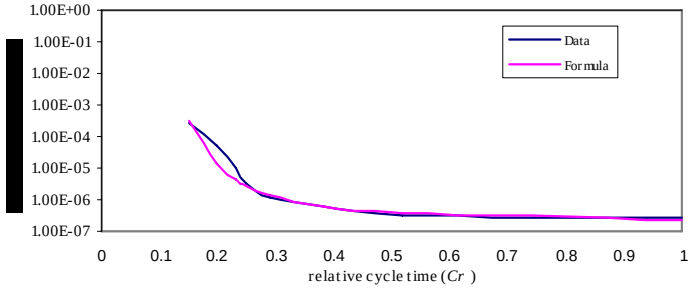

Fault probability versus relative clock frequency

$$P_E = 2.59 * 10^{-7} * e^{\frac{1}{6*C_r^2}} = 2.59 * 10^{-7} * e^{\frac{F_r^2}{6}}$$

Task 2: Measuring Impact on a Processor

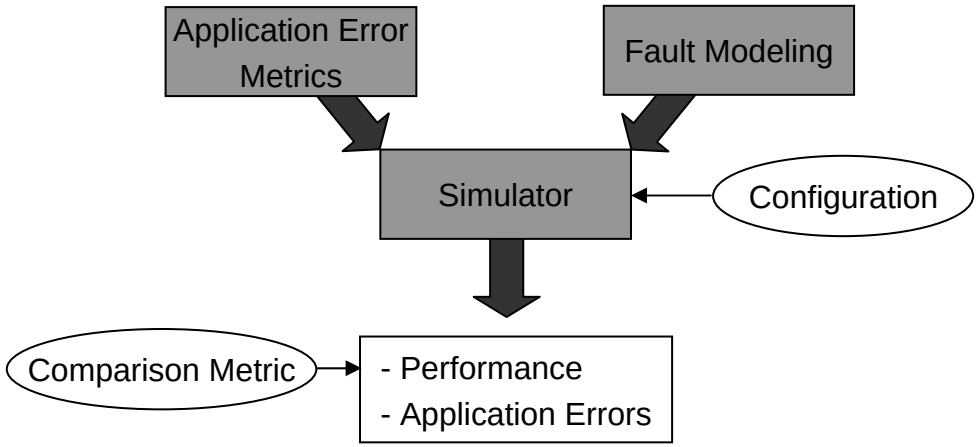

- Mark data structures in apps
  - Important Data Structures
  - Outputs of Key Function Units
- Perform simulation
  - Introduce hardware faults
- Mark the change
- Define the application error rate

Preliminary Results

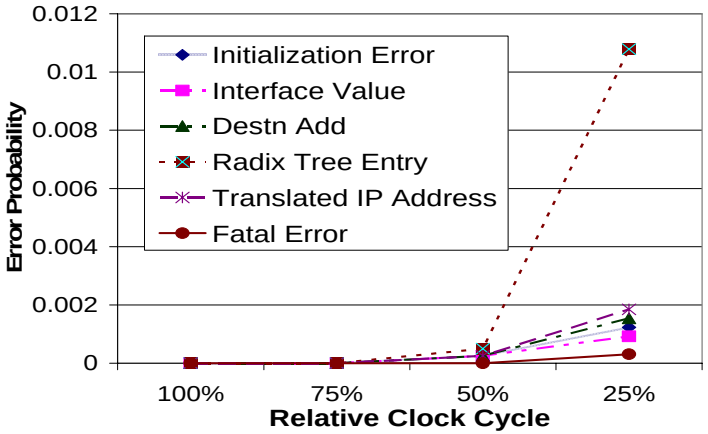

Control plane

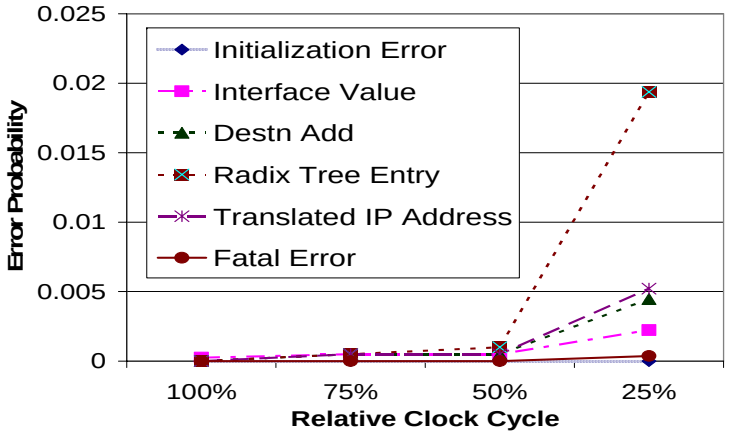

Data plane

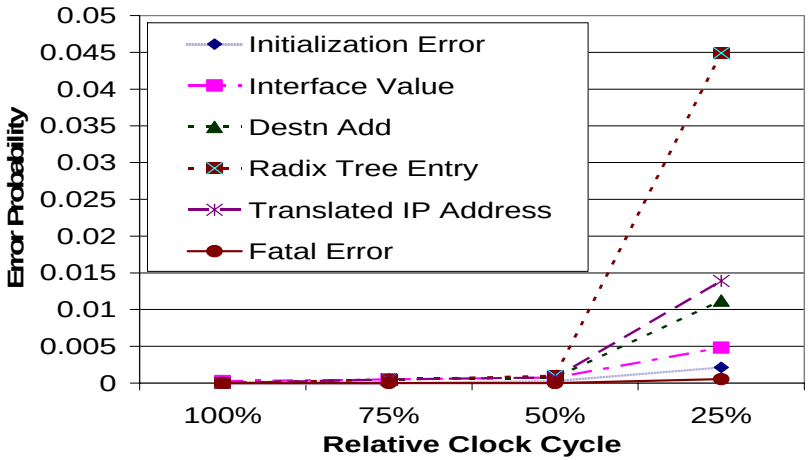

Error introduced in both control and data plane

NORTHWESTERN UNIVERSITY

Activity Plan

- Task 3: System/Network modeling
- Continue on feasibility studies
  - Hardware
    - Simulation framework, testbed, application measurements
- Task 1: New design methodologies
- Task 2: New applications
  - Security

NORTHWESTERN UNIVERSITY

Activities

- Publications
  - A. Mallik, M. C. Wildrick, G. Memik, “Application-Level Error Measurements for Network Processors”, to appear in *IEICE Transactions on Information and Systems*
  - G. Memik, M. Chowdhury, A. Mallik, Y. Ismail, “Engineering Over-Clocking: Reliability-Performance Trade-Offs for High-Performance Register Files”, in Proc. of IEEE/ACM *International Conference on Dependable Systems and Networks* (DSN),Yokohoma, Japan, June - July 2005
  - A. Mallik and G. Memik, “A Case for Clumsy Packet Processors”, in Proc. of *IEEE/ACM International Symposium on Microarchitecture*, Portland, OR, Dec. 2004
  - A. Mallik, M. C. Wildrick, G. Memik, “Measuring Application Error Rates for Network Processors”, in Proc. of *IEEE International Midwest Symposium on Circuits and Systems* (MWSCAS), Hiroshima, Japan, July 2004
- Invited Talks
  - Providing Application-Specific Reliability at the Microarchitecture Level (University of Wisconsin, Apr. 2005)
  - Application-Specific Microarchitecture Optimizations (University of Toronto, Feb. 2005)

Acknowledgements

- We thank the generous support of this project by DOE ECPI program
- Arindam Mallik, Dept. of EECS, Northwestern (PhD student)
- Matthew C. Wildrick, Dept. of EECS, Northwestern (undergraduate)
- Prof. Yehea Ismail, Northwestern

www.ece.northwestern.edu/~memik

memik@ece.northwestern.edu

Thanks!