Tangle: Route-Oriented Dynamic Voltage Minimization for Variation-Afflicted, Energy-Efficient On-Chip Networks

Amin Ansari†, Asit Mishra‡, Jianping Xu‡, and Josep Torrellas†

†University of Illinois at Urbana-Champaign
‡Intel Corporation

HPCA-2014 **Feb 18, 2014**

Trade-off between Energy and Reliability

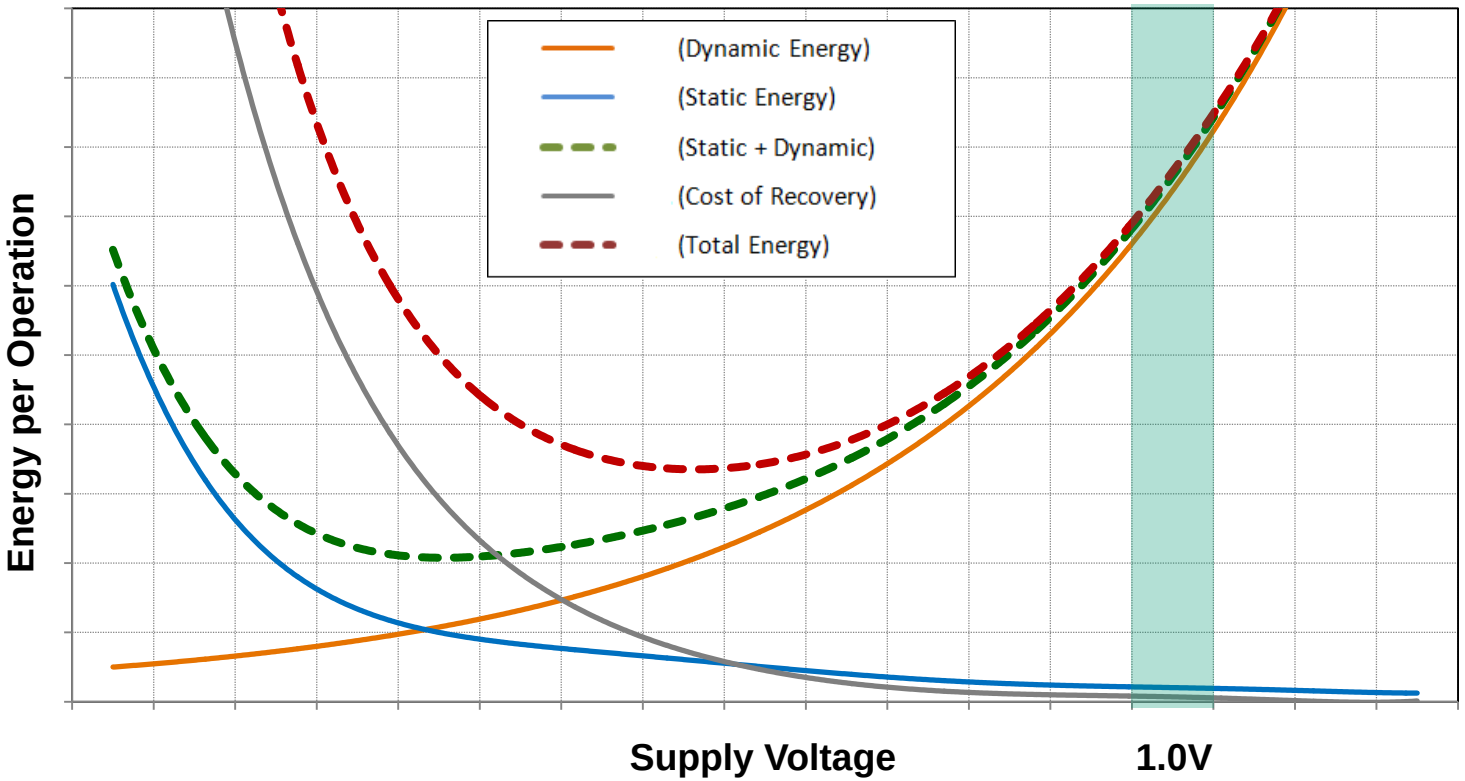

Low-Voltage Operation

- V_{dd} reduction is one of the best levers for energy efficiency
  - Big reduction in dynamic power; also reduction in static power
- New technologies with smaller feature sizes
  - Increase in gate delay variation
  - Requires larger Vdd guard bands to avoid timing faults
- Opportunity: Reduce the conservative Vdd guard-bands while tolerating variations

Motivation: NoC Energy Efficiency

- NoC can consume up to 40% of chip power [Li'08]
- Especially vulnerable to variations
  - Routers and links connect distant parts of a chip
  - Exhibit different speed and power characteristics
- Focus of this paper:
  - **Reduce V_{dd} guard-bands in NoCs to save energy**

Contributions

- Proposed a mechanism (**Tangle**) to reduce the Vdd guard-bands in NoC routers affected by variation

- **Tangle approach**
  - Organize the NoC into multiple Vdd domains
  - Periodically lower the Vdd of the each domain
  - Monitor any message errors
  - On error, increase the Vdd of the routers in the message path
  - Each domain converges to its lowest error-free Vdd

- **Results**: Energy reduction in a 64-node NoC:
  - 28% with 1 router/domain and 22% with 16 routers/domain
  - Negligible performance impact

Fmax/Leakage Var. in Intel's SCC

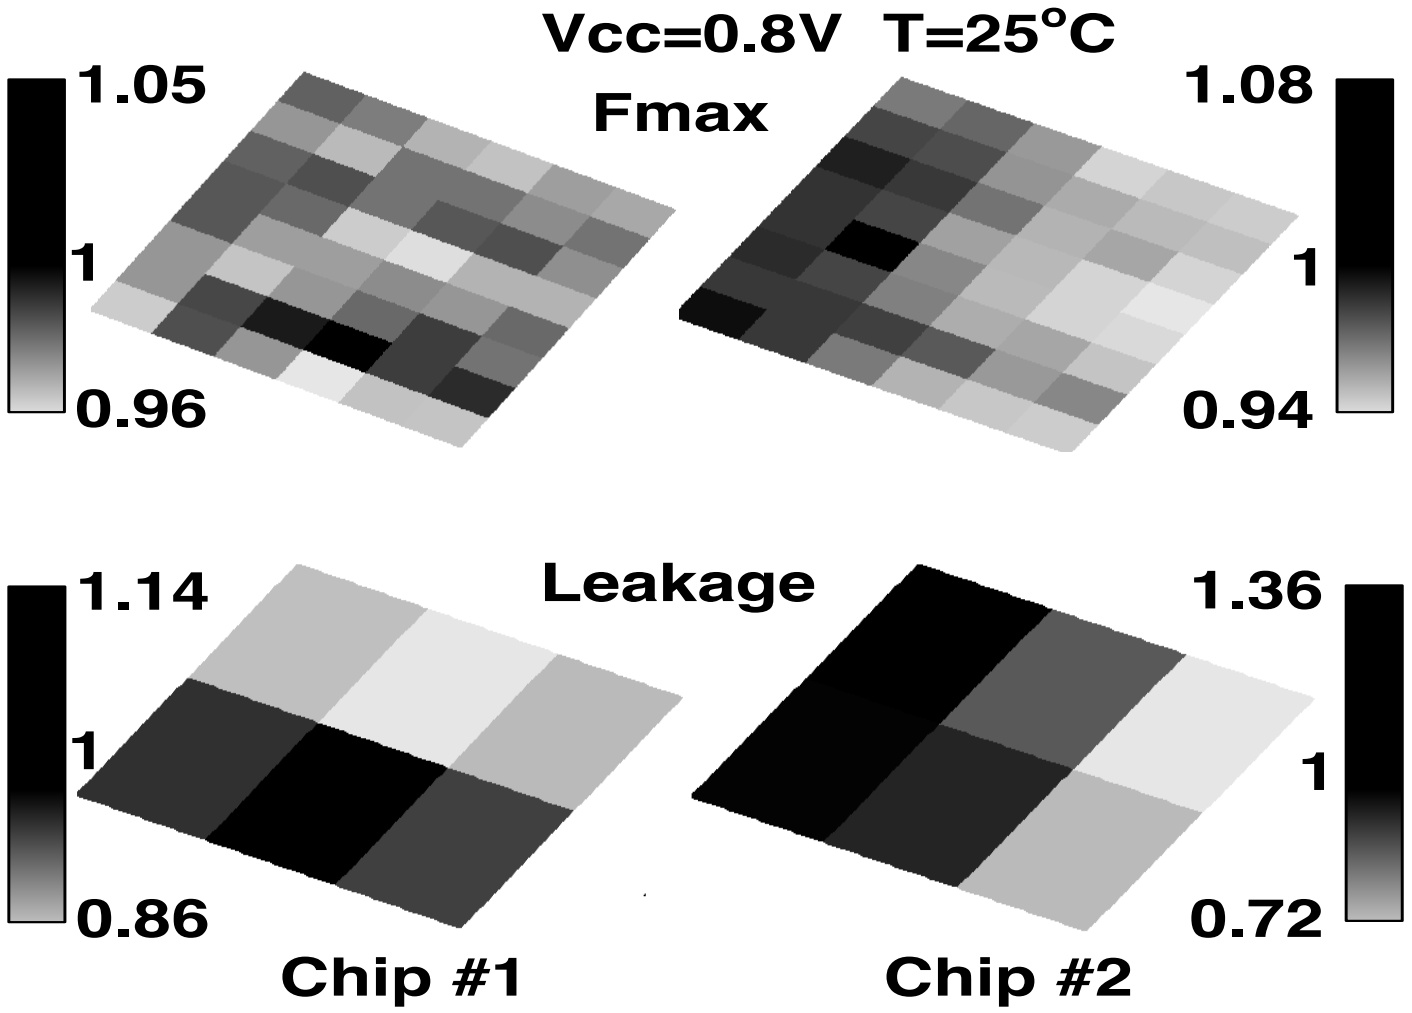

Error Rate as Function of V_{dd}

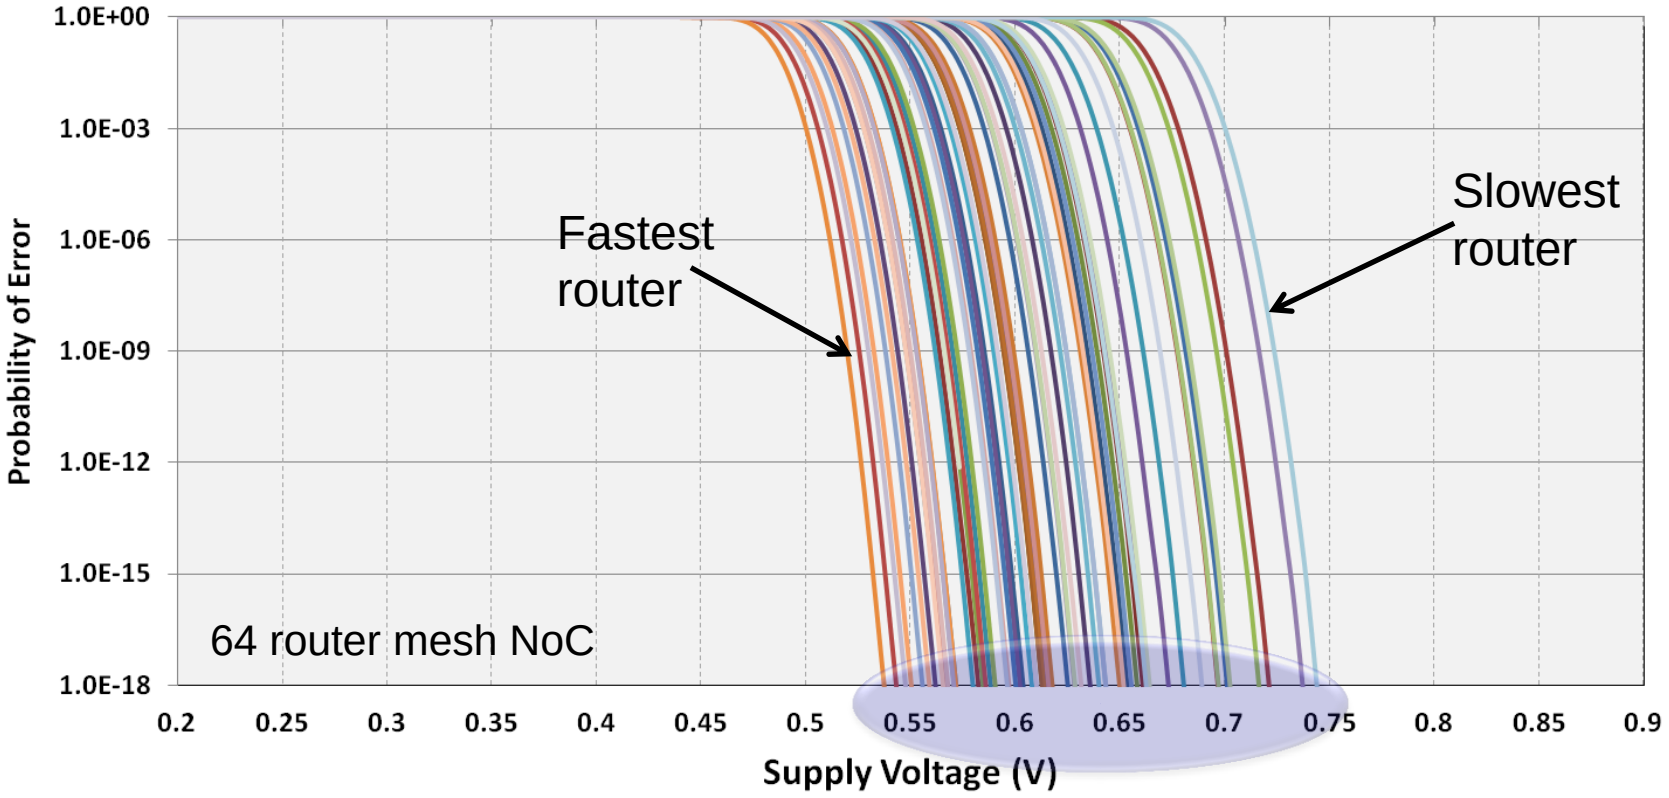

Process variation has a major impact on the routers

Tangle: Key Idea

- Attain high energy efficiency by minimizing V_{dd} margins added for variation and wearout
  - Constant frequency
- Reduce V_{dd} for each group of routers
  - Starting at a high voltage (e.g., 800mV)
  - Decrease V_{dd} periodically
  - While continuously monitoring for errors
    - Inexpensive error detection: CRC end-to-end

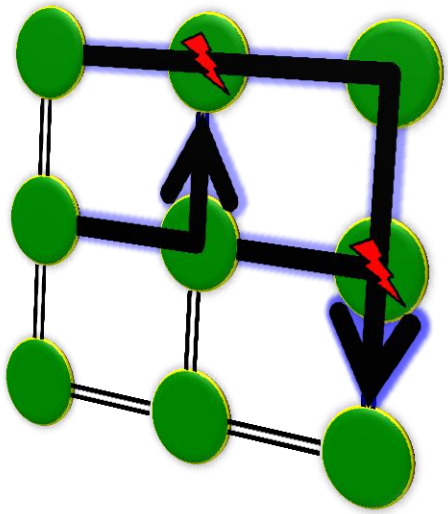

- Dynamically change V_{dd} of each group of routers
  - Adapting to errors seen during changes in workload phases, temperature, and also wearout

Tangle: Handling Errors

Continuously monitors errors, when an error occurs

- Destination node drops the flit and waits for a retransmission
- At the source node, we have a watchdog timer. When watchdog timer times-out:
  (1) Source sends a signal to Reliability Management Unit (RMU) for $\uparrow V_{dd}$
  (2) Source retransmits the message

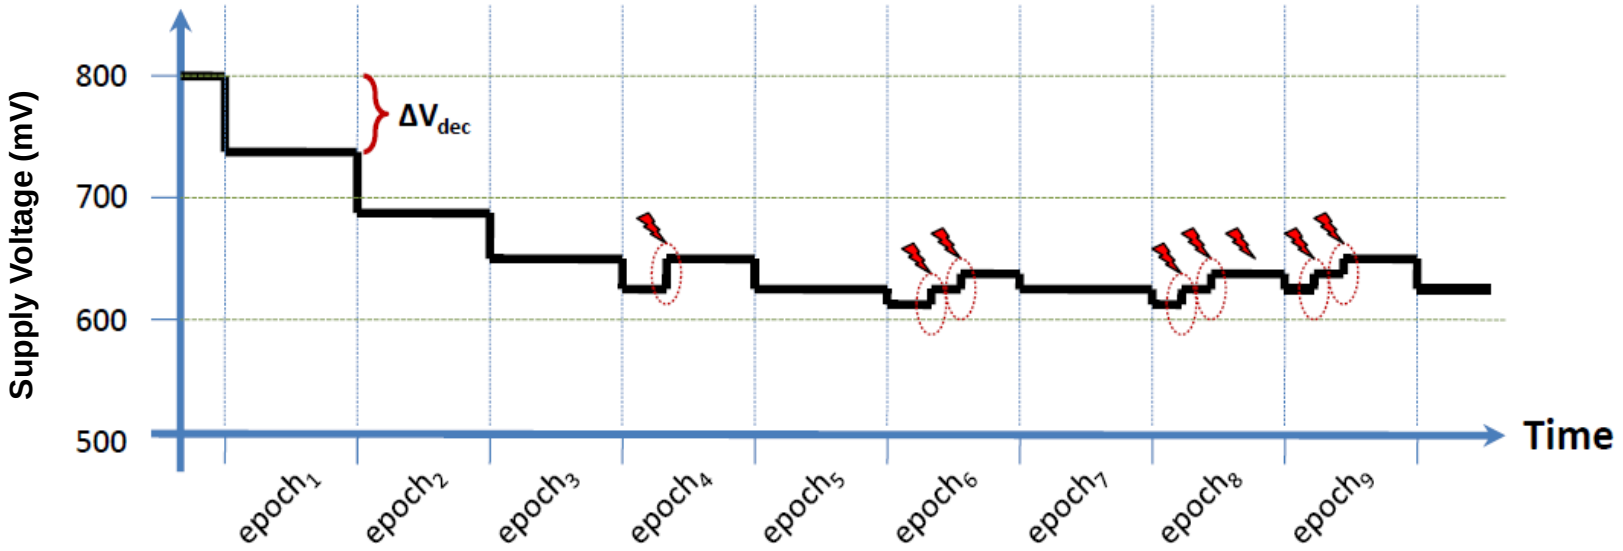

Error Detection

- We want a low-cost error detection
  - No significant increase in the traffic
    - Only a few bits should be added to every message
    - Error checking should be done infrequently
  - Error detection code needs to be strong
    - We can have multiple bit errors in a message (bursts)
    - Hamming, parity, and simple detection codes are not sufficient
  - We chose CRC
    - Apply it as an E2E error detection mechanism
    - Pick WCMDA-8 which is close to a perfect 8-bit CRC
    - Years of operation without error

Initial State of the System

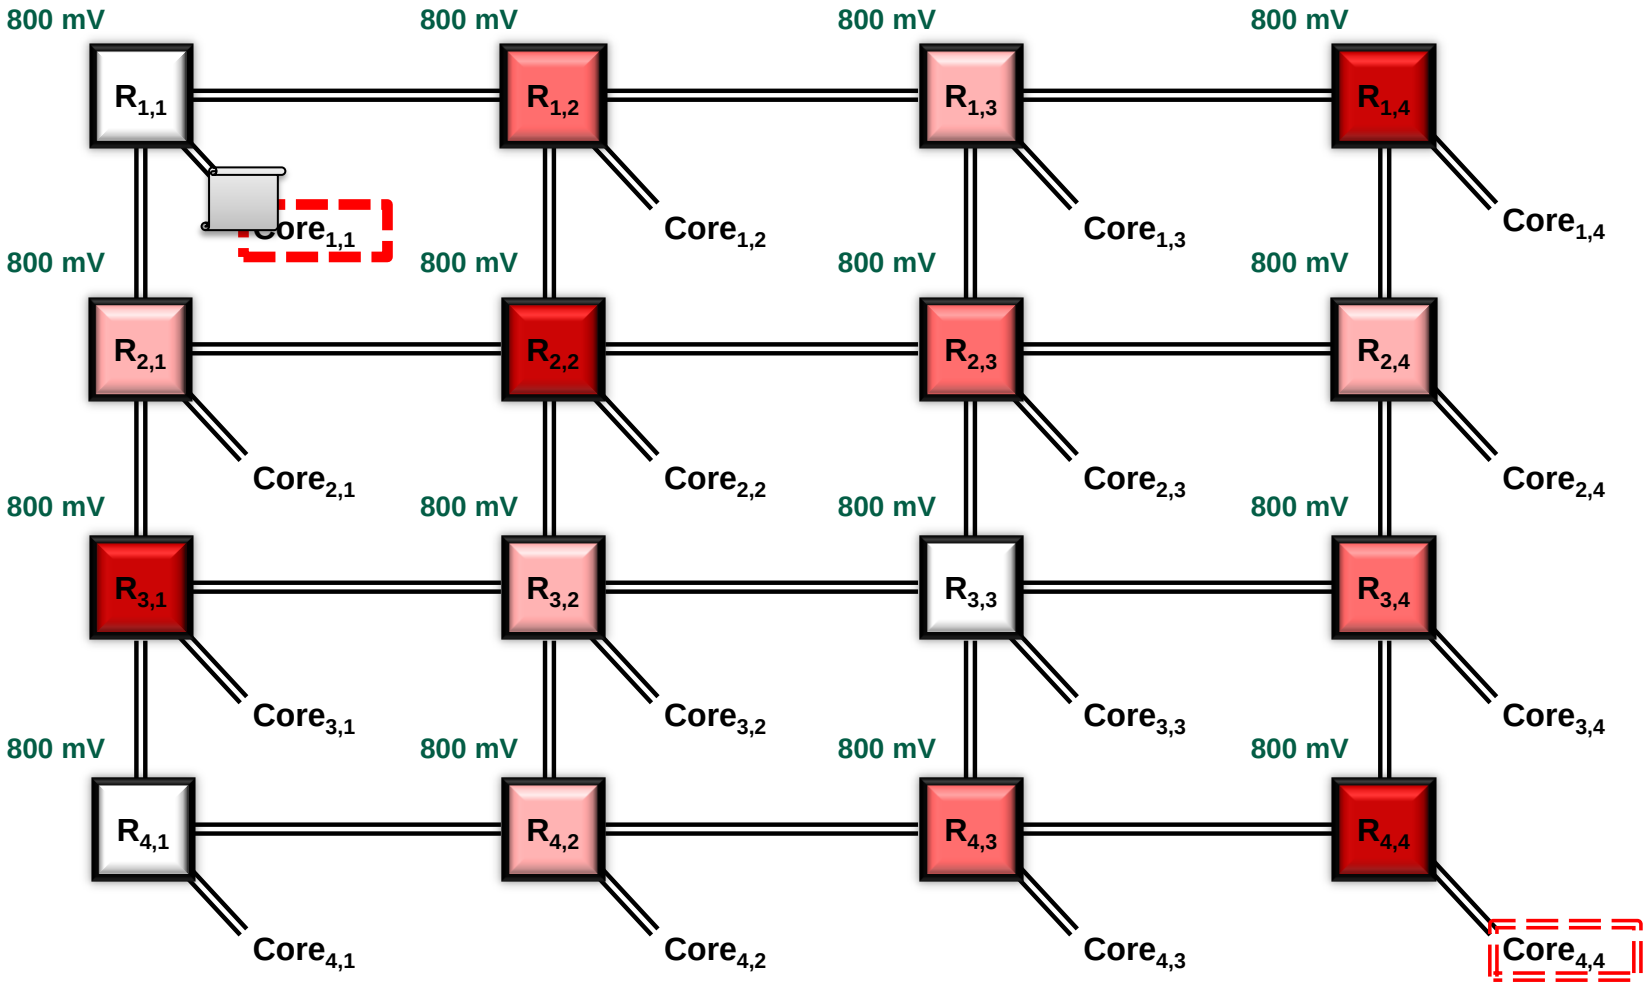

System After a Few Epochs

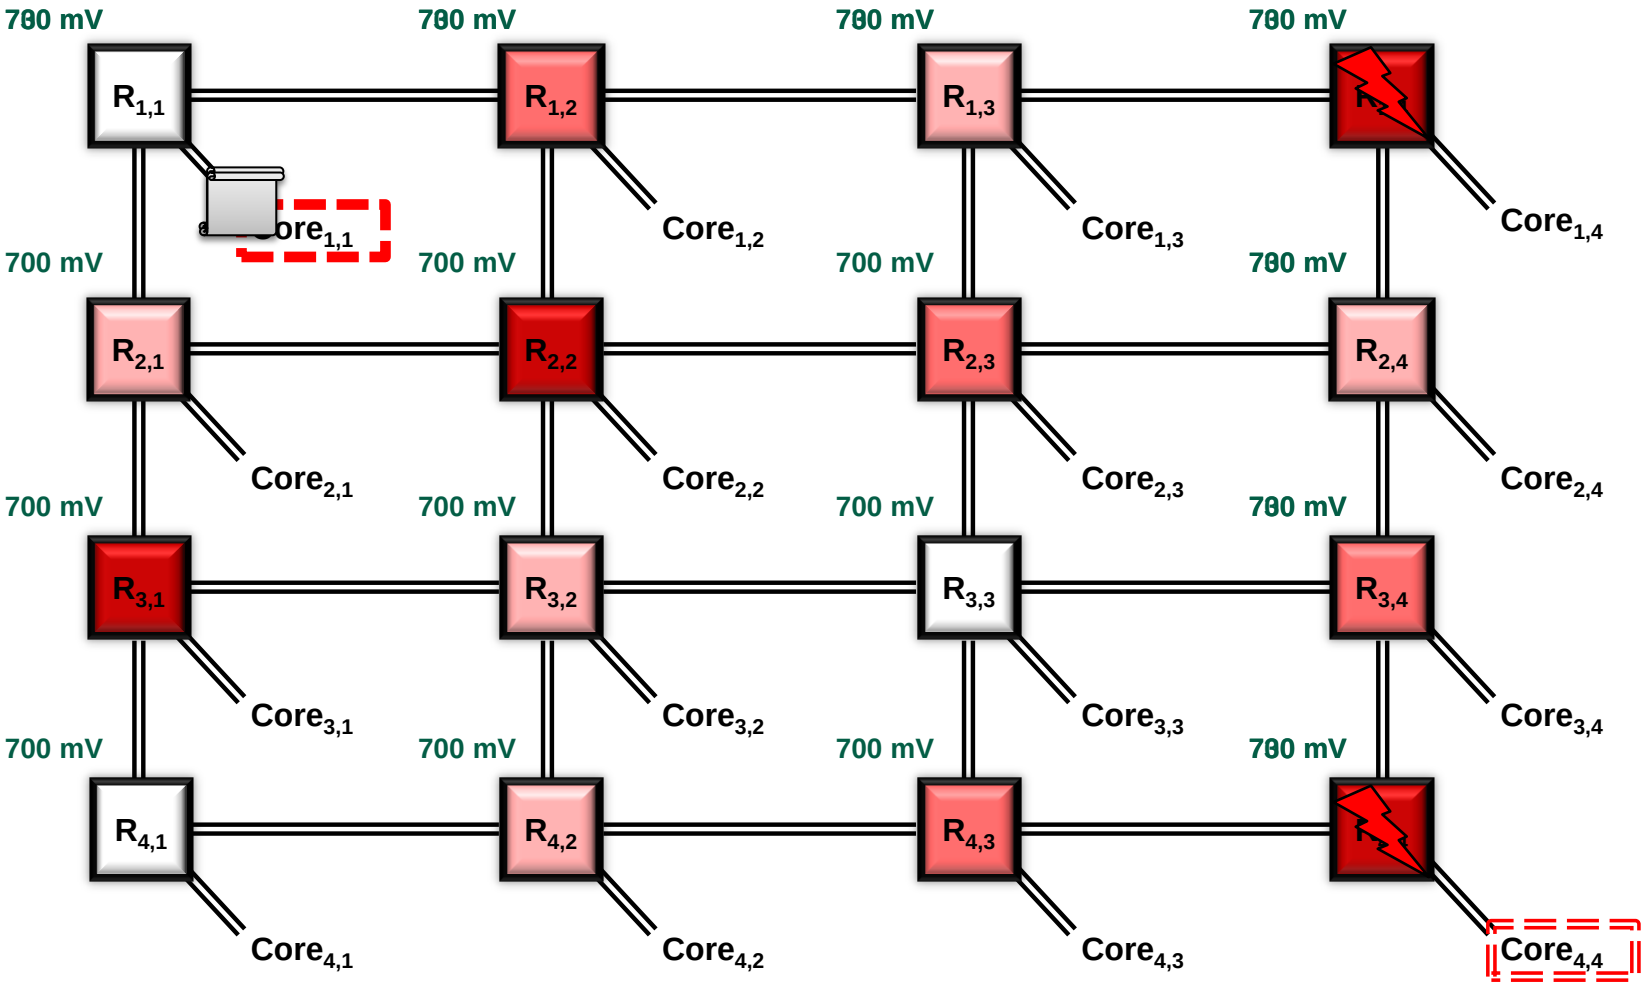

Steady State after Enough Epochs

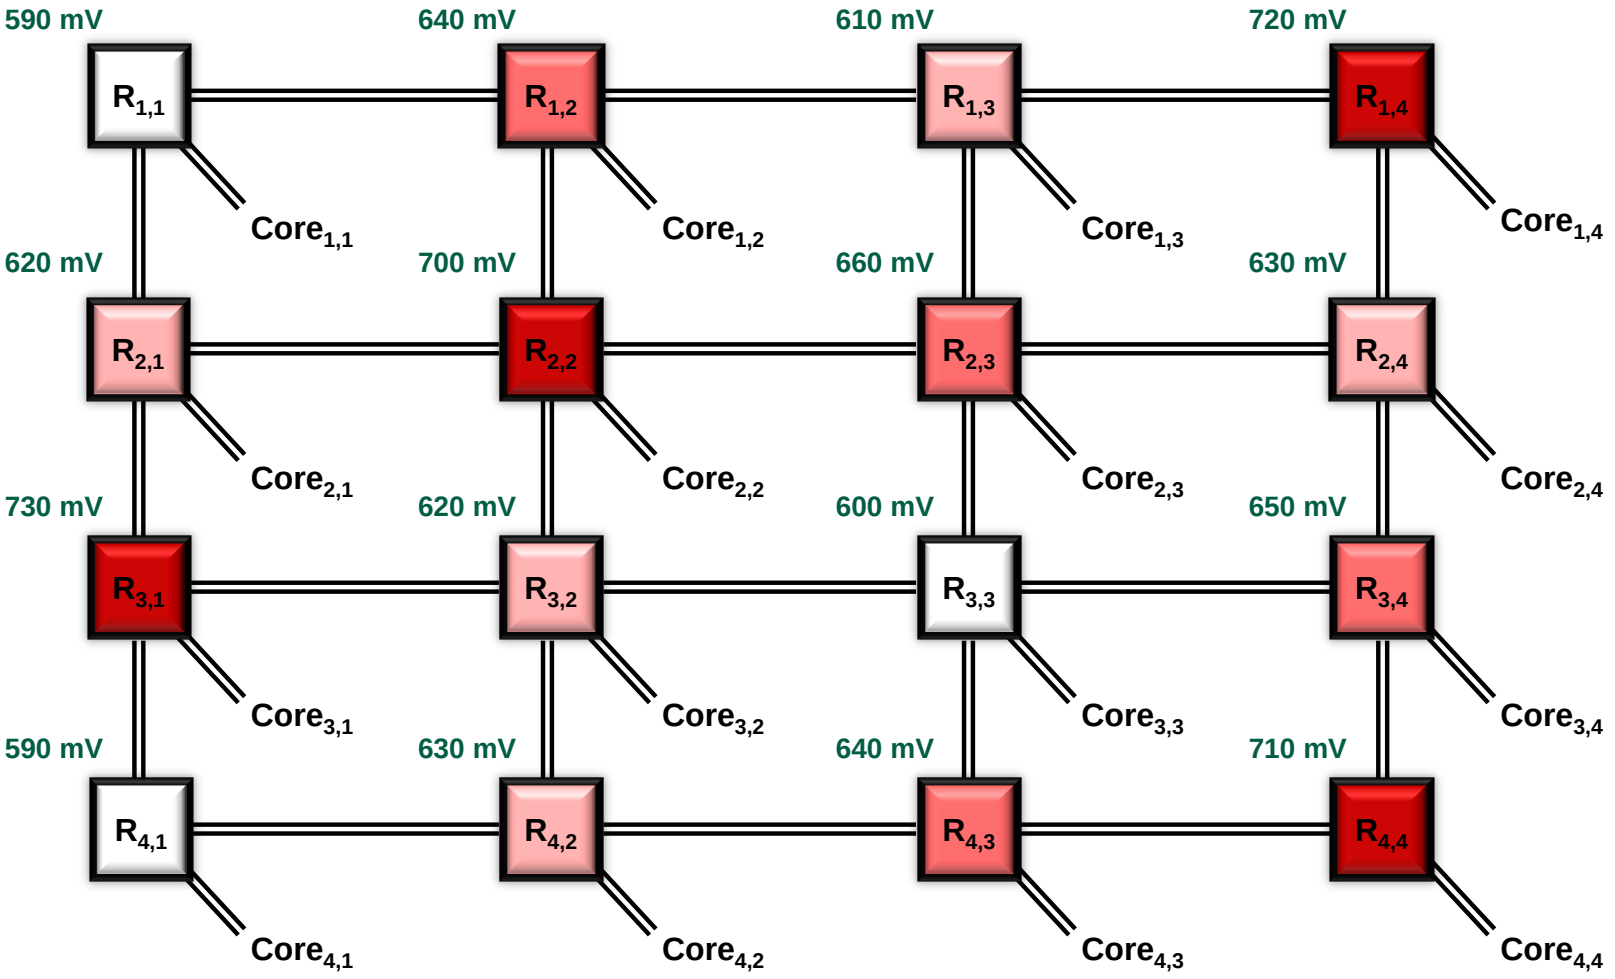

V_{dd} Tuning

- Performed by Reliability Management Unit (RMU)
  - Parallel low-bandwidth network
  - Manages the V_{dd} of routers in each group (voltage domain)
  - Computes path from sender to receiver

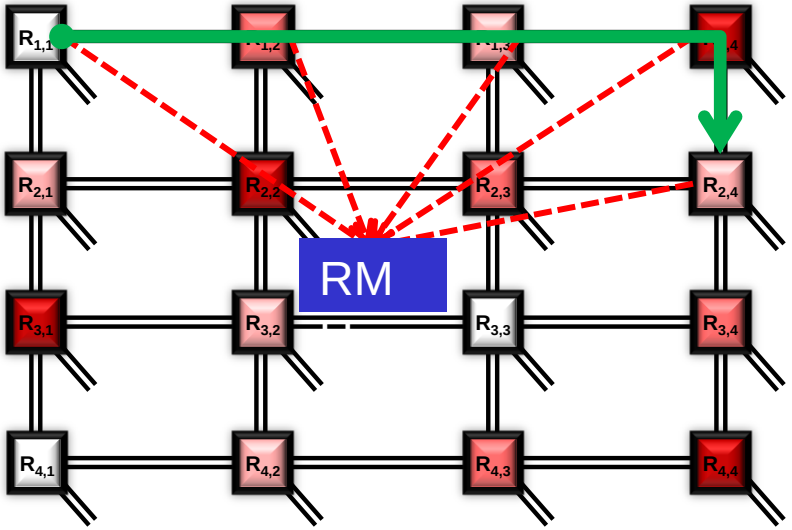

V_{dd} Tuning

- Performed by Reliability Management Unit (RMU)
  - Hierarchically built (H-tree structure)
  - Voltage tuning done in a gradual manner (high Δ values in initial epochs followed by low Δ values in later epochs)
  - Only allows finite tuning steps in an epoch.

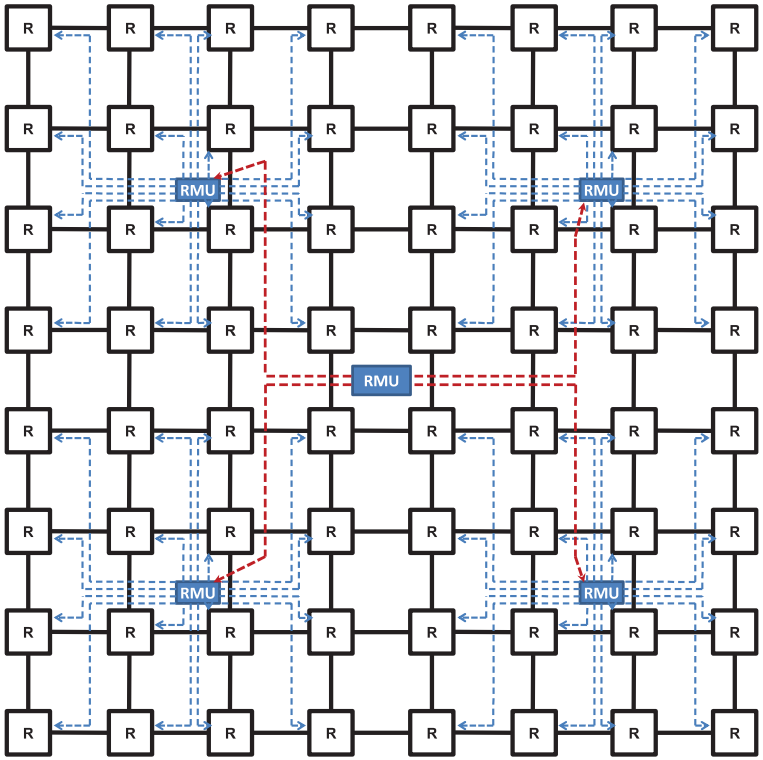

Costs of Tangle Design

- Voltage regulation
  - Support for multiple Vdd domains
  - Loss of ~10% power
- Error checking
  - 8b CRC added to each 128b flit
  - Negligible performance overhead
- RMU and its narrow-network
  - Very small area and power
  - Off the critical path
- Message re-transmission
  - Happens relatively rarely
  - Small overhead

Variation Modeling

- RTL model of a three stage router
  - Modified Verilog model from Peh'01
- Synopsys design compiler
  - 45nm Nangate open cell-library
- Selected 32 close to critical paths for analysis per stage
  - Synthesis static timing results
- Apply variation for these paths
  - VARIUS for process variation modeling and error rates
  - Error rates generated for given voltages and frequencies
  - Modeling random and systematic variation

Microarchitectural Simulation

- Multicore
  - Homogenous CMP
  - Dual issue OoO Alpha processors
  - Distributed L2 with static address mapping
- Network
  - Mesh network (vary size: 4x4, 6x6, 8x8 (default), 10x10)
  - Deterministic routing and credit-based flow control
  - Vdd domains: fine-grained (1 router/domain) to coarse-grained (16 routers/domain)
- Workloads
  - 3 commercial workloads (sap, sjas, tpcw), 9 benchmarks from SPEC-CPU-2006, 3 scientific benchmarks (art, ocean, and swim)

Steady-State Voltages

730	630	650	700	690	680	620	720
630	670	600	620	660	680	600	590
660	650	680	660	670	640	660	660
700	590	610	630	680	610	630	690
680	640	660	610	640	600	690	730
640	650	640	620	670	610	630	660
630	620	600	610	710	630	650	630
700	660	670	670	670	680	640	660

Nominal V_{dd} = 800mV
SAP workload

Observations:

- **Spatial correlation because of systematic process variation**
- **Wide voltage range from 590mV all the way up to 730mV**

Energy Saving Results

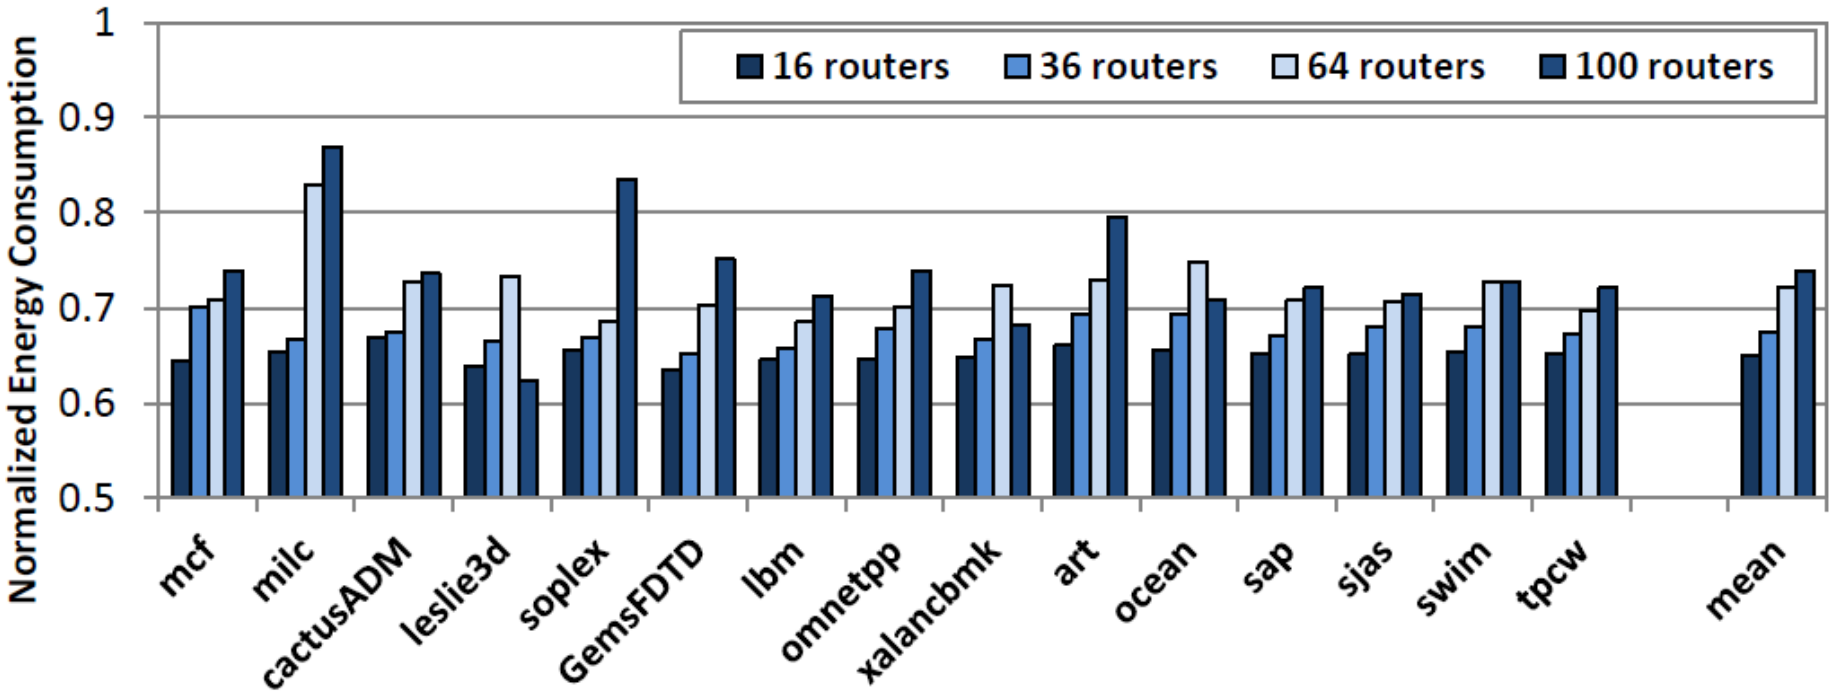

Energy Saving Results

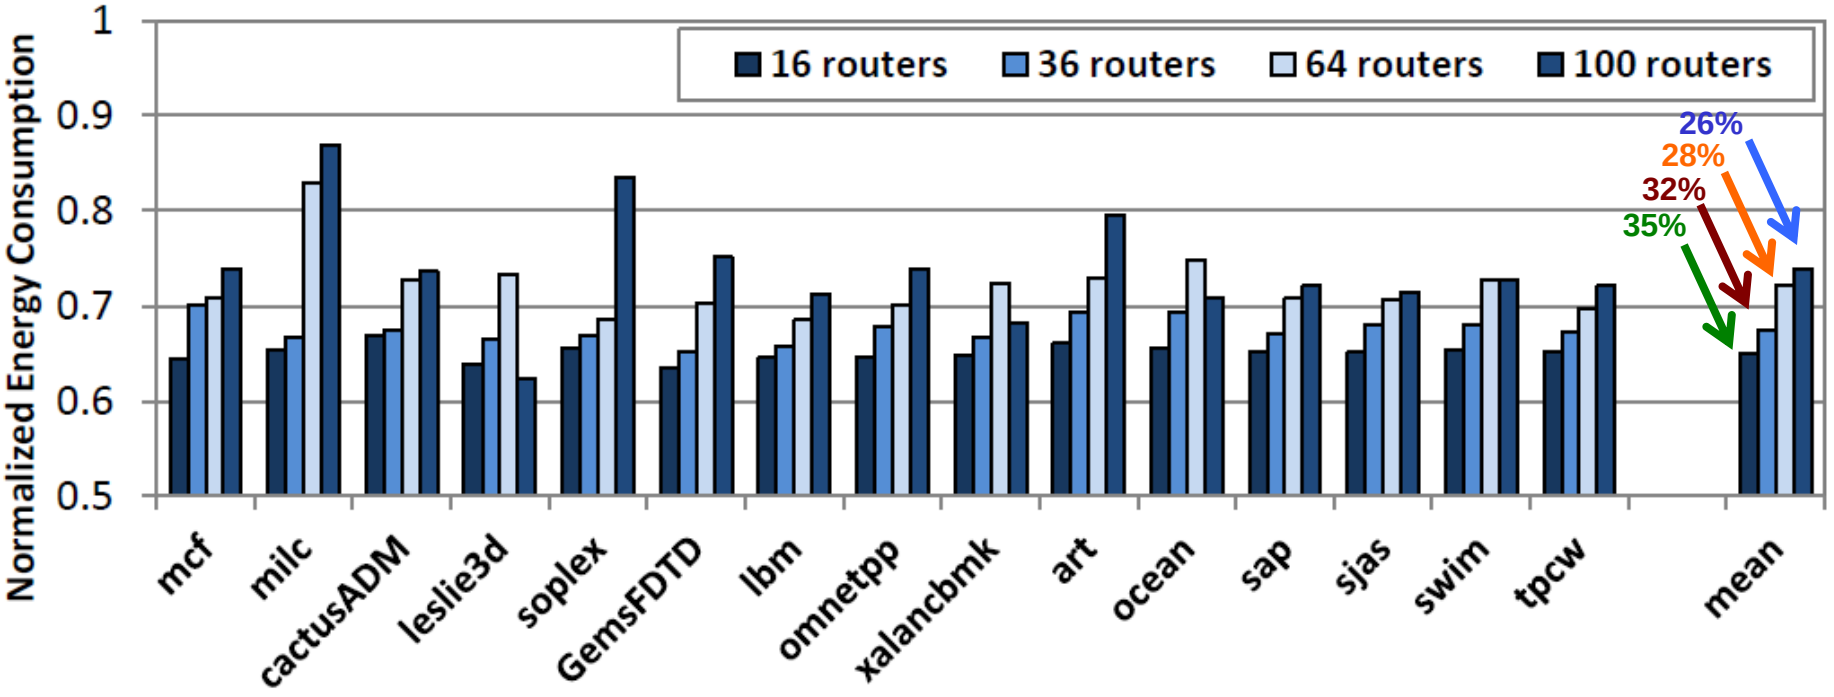

Average energy saved for 16, 36, 64, and 100 router NoCs is 35%, 32%, 28%, and 26%, respectively (with less than 1% performance loss)

Energy Saving Results – contd.

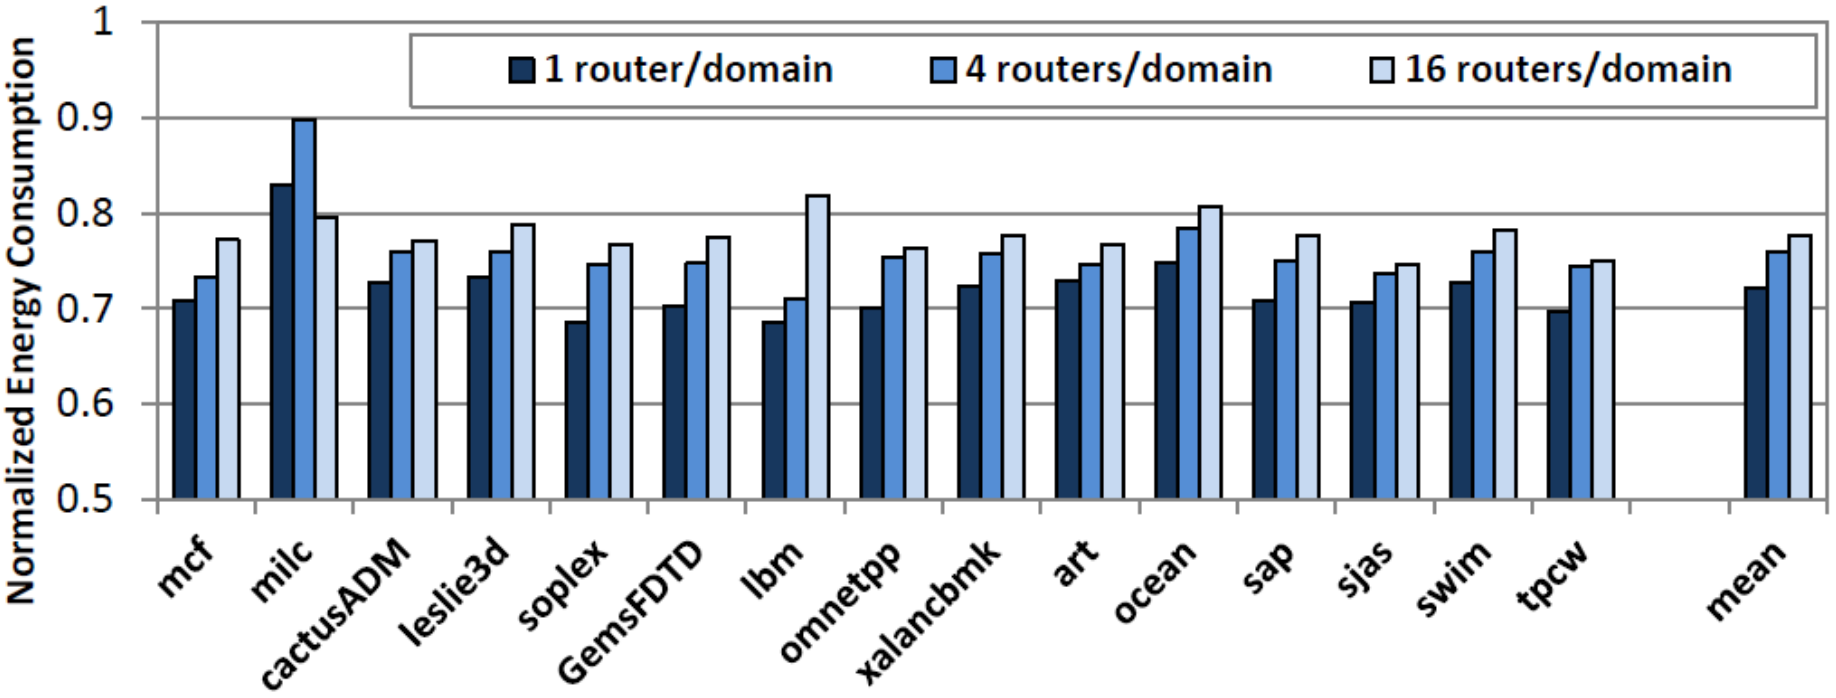

Energy Saving Results – contd.

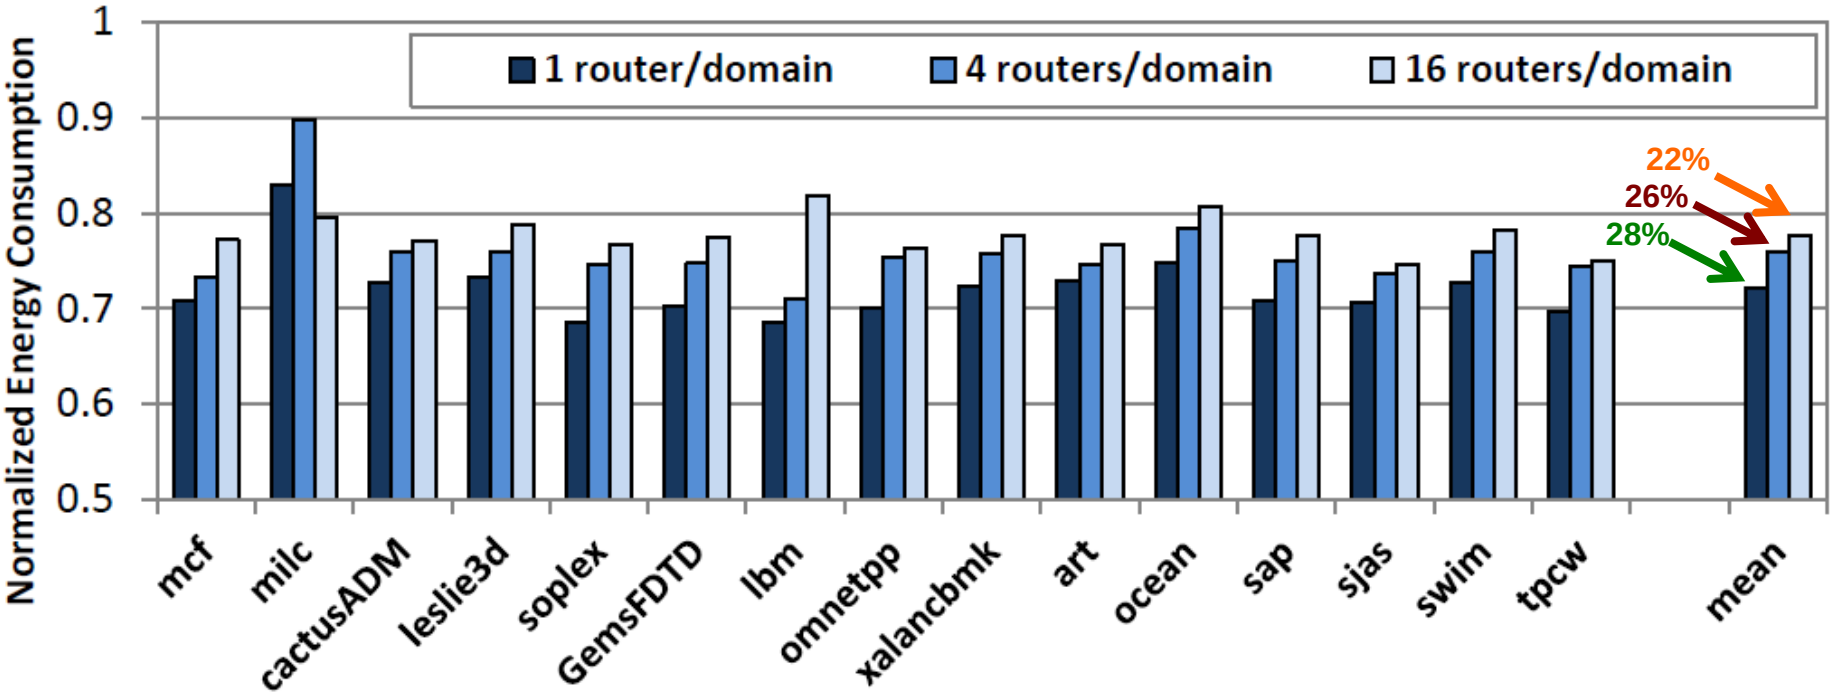

Average energy saved for 64 router NoC is
28% (with 1 router/domain), 26% (4 router/domain) and 22% (16 router/domain)

Conclusion

- Proposed a mechanism (**Tangle**) to reduce the Vdd guard-bands in NoC routers to save energy
  - In the presence of process variation
- Tangle approach
  - Monitor the message errors
  - Increase the Vdd of the router domains in the message path
- Tangle applied to a 64-node mesh network
  - Reduces network energy by 28% (avg.) with 1 router/domain and 22% with 16 routers/domain
  - Negligible performance impact

Thank You