

Wear-Aware Adaptive Routing for Networks-on-Chips

Arseniy Vitkovskiy[†], Vassos Soteriou[†], Paul V. Gratz[‡]

[†]Dept. of Electrical Eng., Computer Eng. and Informatics
Cyprus University of Technology
vassos.soteriou@cut.ac.cy

[‡] Dept. of Electrical and Computer Engineering
Texas A&M University
pgratz@tamu.edu

ABSTRACT

Chip-multiprocessors are facing worsening reliability due to prolonged operational stresses, with their tile-interconnecting Network-on-Chip (NoC) being especially vulnerable to wearout-induced failure. To tackle this ominous threat we present a novel wear-aware routing algorithm that continuously considers the stresses the NoC experiences at runtime, along with temperature and fabrication process variation metrics, steering traffic away from locations that are most prone to Electromigration (EM)- and Hot-Carrier Injection (HCI)-induced wear. Under realistic applications our wear-aware algorithm yields 66% and 8% average increases in mean-time-to-failure for EM and HCI, respectively.

Categories and Subject Descriptors

B.4.3 [Input/Output and Data Communications]: Interconnections (Subsystems); B.8.1 [Performance and Reliability]: Reliability, Testing, and Fault-Tolerance

General Terms

Reliability, Design

Keywords

Hot-Carrier Injection (HCI), Electromigration (EM), Network-on-Chip, Routing Algorithm, Reliability, Lifetime

1 Introduction

Today's Chip Multiprocessors (CMPs) utilize tens or hundreds of processing cores that coordinate to execute applications in parallel. In these systems, processors, memories, IP blocks and peripherals, are serviced by high-throughput Networks-on-Chips (NoCs). Unfortunately, operational stresses imposed by CMP workloads make such chips increasingly prone to accelerated wearout, with recent studies projecting a 10× increase in wear-induced failure over the next 10 years [4]. The NoC is especially vulnerable to such failure, as a *single* component malfunction could render the CMP unusable [4]. NoC failure is unlike individual core failure, where error detection and appropriate operating system support can be used to re-map executing tasks to healthy cores [3].

Much prior work examines fault-tolerant routing, attempting to *reactively* manage faults as they occur [5]. Ideally, *proactive* schemes could extend the healthy status of the system without incurring faults, rather than to react to such fault occurrences. Thus, in this paper we develop a proactive *Wear-Aware Routing Algorithm* (WARA), which restricts network traffic-induced wear in the NoC. Our algorithm considers the stresses that NoC routers experience online, focusing on the data- and activity-dependent wearout effects of Hot-Carrier Injection (HCI) and Electromigration (EM) that are most critical to the interconnect. Using micro-architectural-level models for HCI and EM wear, along with dynamic operating temperature, and process variation, WARA steers traffic in the NoC topology

to reduce wear on the NoC. WARA avoids those NoC components or routing paths most susceptible to wearout due to process variation, higher operating temperatures, or accumulated wear to-date. Under applications from the PARSEC Benchmark Suite, WARA yields an 66% and 8% increase in mean-time-to-failure for EM and HCI, respectively.

2 Background on CMOS Physical Wearout

Two primary CMOS failure mechanisms are EM [1] and HCI [2]. In both cases, the proximate cause of failure is the activity factor (probability of bit transition) of the wire (EM) or transistor (HCI) under stress. Wear rate is further modulated by the temperature (in both EM and HCI), and process variation (PV) (in HCI). We next present models, for temperature and PV, as these are the primary aging accelerators for EM and HCI.

2.1 Dynamic Temperature Variation

Chip operating temperature has a direct influence on both EM- and HCI-induced aging. CMPs exhibit uneven power dissipation among their component resources, dependent upon the specific task mapping, application load balance, etc., all of which are inherently unbalanced. We assume a per-tile random temperature distribution driven by the core-produced heat where the temperature remains constant within a tile but varies between tiles. The temperature range spans from 65 °C to 85 °C, with values altered periodically to mimic the CMP's dynamic behavior (see Section 3).

2.2 Process Variation (PV)

Process variation (PV) is a deviation in technology process parameters from their nominal values caused by imprecision in the fabrication process as feature size approaches the fundamental dimensions, leading to imbalance of electrical characteristics across the die. Within-die PV is a combination of *systematic effects* (e.g., layout pattern density variation), and *random effects*. Since variation has a very small effect on wires [7], we consider the process parameters most vulnerable to variation: the transistor threshold voltage, V_{th} , and the effective gate length, L_{eff} . The values of the systematic components, $\Delta V_{th,sys}$ and $\Delta L_{eff,sys}$, are assumed to be constant within one CMP tile. These may be modeled using a multivariate normal distribution with a spherical spatial correlation between the CMP tiles. The finer-grained (i.e., at the level of an individual transistor) random components, $\Delta V_{th,rand}$ and $\Delta L_{eff,rand}$, are assumed both normally distributed with zero mean and standard deviation equal to 6.3% and 3.2%, respectively [7].

2.3 Electromigration (EM)

EM affects metallic wiring on-chip, where current-driven stresses, as a result of bit transition switching, causes deformation due to the movement of metal ions, leading to fatal opens or shorts. Of particular concern for EM are the Vdd and GND connections of drivers for high-rate data transfer wires, where both current drive is high, due to long wire capacitance, and switching activity is high, exacerbated by high temperature. The *Acceleration Factor* (AF) is a relative metric of lifetime [2] for EM wear. Here, we assume the reference functional lifetime as that of a network with deterministic Dimension-Order Routing (DOR), where the EM

acceleration factor is calculated as $AF_{EM} = \frac{MTTF_{EM}|_{AR}}{MTTF_{EM}|_{DOR}}$,

where $MTTF_{EM}|_{AR}$ is the Mean-Time-To-Failure (MTTF) of the system, which runs WARA, and $MTTF_{EM}|_{DOR}$ is the reference MTTF of the standard DOR-running system.

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the Owner/Author.

Copyright is held by the owner/author(s).

NOCS '15, September 28-30, 2015, Vancouver, BC, Canada

ACM 978-1-4503-3396-2/15/09.

<http://dx.doi.org/10.1145/2786572.2786573>

2.4 Hot-Carrier Injection (HCI)

HCI affects transistors in proportion to the time they have been operating under stress, *i.e.*, the switching frequency. HCI causes gradual shifts of a transistor's threshold voltage (V_{th}), which slows its switching time, eventually leading to timing guard-band violation, affecting a device's critical path. HCI has a further sensitive to process variation (PV) as PV can cause an initial shift in V_{th} , reducing the guard-band before HCI has even begun to effect the transistor. The Acceleration Factor (AF) for HCI is calculated as

$$AF_{HCLPV} = \frac{MTTF_{HCLPV}|_{AR}}{MTTF_{HCLPV}|_{DOR}}, \text{ where } MTTF_{HCLPV}|_{AR}$$

is the MTTF of the system, which runs wearout-aware adaptive routing, and $MTTF_{HCLPV}|_{DOR}$ is the reference MTTF of the system, which runs standard DOR.

3 Wear-Aware Adaptive Routing Algorithm

As discussed above, the cause of aging under both EM and HCI wear-inducing physical phenomena is switching activity, caused by packets being routed. Load within an NoC is often quite unbalanced, thus reducing the switching activity of nodes and links which have the greatest switching activity, can be done by using adaptive lifetime-extending routing that does not compromise the utility of the network.

Our proposed WARA leverages a holistic view of the network's wear-sensitivity, conveying distant accumulated wear to-date, current temperature state, as well as initial PV state, to the local router for the creation of a wear-reducing route. In gathering global statistics, WARA inherits features from a congestion-avoiding, globally-aware adaptive routing scheme, *Global-Congestion Awareness (GCA)* [6]. It holistically communicates wearout status information by "piggy-backing" data that is back-annotated into the empty space of header flits comprising (1) utilization to-date of the link that resides in the opposite direction of the packet's flow (U_{link}), (2) utilization to-date of the router's critical path (U_{router}), and (3) temperature of the corresponding network tile. The U_{link} and U_{router} statistics are used to continuously estimate both EM and HCI activity-based wear, respectively, implemented as 32-bit counters, one for each router port, that are incremented with every flit's link traversal. At each hop, every node appends this information into the header flit and then sends it out on its associated output port. Thus information, sampled every 10^4 cycles, from one router spans the entire topology and is received by all other routers. The average U_{router} metric is then estimated by summing all per-port counters and dividing by the packet's flit count. Last, fabrication-time PV constants are measured and stored once in tables at each router, while temperature data for each tile, also stored in router tables, are sampled every 10^7 cycles.

Activity statistics, PV and temperature values, are all fed into the lifetime-determining equations of Black [1] and Hoskote *et al.* [2] that respectively calculate the MTTF of electronic devices for EM- and HCI-induced wear. WARA's goal is to choose the packet route among the alternative progressive paths that incurs the least cumulative wear, for both EM and HCI in tandem. The path-accumulated, relatively-weighted lifetime, W_{norm} , from a given source node to a final destination, out of the set of all in-between minimal path permutations is proposed as WARA's routing function:

$$W_{norm} = \sum_{minpaths} (C(MTTF)_{EM} \times MTTF_{EM,norm_i} + C(MTTF)_{HCLPV} \times MTTF_{HCLPV,norm_i}) \quad (1)$$

where

$$C(MTTF)_{EM,HCLPV} = \frac{exp^{w \frac{MTTF - MTTF_{max}}{MTTF_{max} - MTTF_{min}}} - exp^{-w}}{1 - exp^{-w}} \quad (2)$$

where $MTTF_{EM,norm_i}$ and $MTTF_{HCLPV,norm_i}$ are normalized to the maximum MTTF values observed in the *least worn-out* path. $C(MTTF)_{EM}$ and $C(MTTF)_{HCLPV}$ are exponential functions that are specifically tuned to apply much lower weights to paths which are close to the lowest

MTTF values, such that they are much less often chosen by WARA. $MTTF_{max}$ and $MTTF_{min}$ denote the maximum and minimum MTTF values along all alternative progressive paths residing between a source-destination pair. The exponential weight constant w scales the strength of the extra weight away from paths with small MTTFs, and is empirically set to 8. The intuition here, is that the path with the least relative MTTF value can be regarded as the "*weakest link*" in the topological "chain" - it determines the NoC's overall lifetime duration and thus needs to be preserved as long as possible. Hence, WARA guides packets to traverse the path with the highest W_{norm} , helping to unload the paths with shorter cumulative lifetime. Route calculation is implemented via a modified Bellman-Ford [6] shortest path algorithm which chooses the progressive path of greatest relative MTTF value, by utilizing weights obtained from Eq. 1. Hence, packets use previously precomputed routes.

4 Experimental Setup and Evaluation

The experimental platform is a 64-core CMP organized as an 8×8 array of tiles, each comprising an x86 processor core, split L1 and L1d private caches, and a slice of the combined, shared L2 cache. A directory-based MESI protocol maintains cache coherence, and uses 3 virtual networks. We utilize 3-stage pipelined wormhole routers, with 4 virtual channels per router port, and 128-bit links. Experiments are run using the gem5 full-system simulator, with the *Ruby* shared memory model and the *Garnet* network simulator. Our workload comprises the PARSEC Benchmark Suite.

We compare the respective CMP lifetimes achieved under each PARSEC benchmark run using WARA vs. a system with a baseline, *wearout-unaware* DOR routing algorithm. Each simulation is repeated fifteen times with different PV and temperature mappings and the final results are averaged for each value obtained. In our MTTF Acceleration Factor (AF) for EM and HCLPV results, the bottom 10% of observable MTTF cases were considered, since those are the *least reliable* and hence most likely to lead to failure. WARA increases system lifetime under both EM and HCLPV, for all tested PARSEC benchmarks, providing 66% increase in $MTTF_{EM}$ and 8% increase in $MTTF_{HCLPV}$, on average. Finally, WARA experiences a small performance degradation, at 2% on average versus a baseline, DOR design.

5 Conclusions

We presented a novel, dynamic wear-aware routing algorithm for NoCs that continuously focuses on combating HCI- and EM-induced wear in tandem. It yields an average 66% increase in $MTTF_{EM}$ and 8% increase in $MTTF_{HCLPV}$.

Acknowledgements

This work falls under the Cyprus Research Promotion Foundation's Framework Programme for Research, Technological Development and Innovation 2009-10 ($\Delta E \Sigma M H$ 2009-10), co-funded by the Republic of Cyprus and the European Regional Development Fund, and specifically under Grant IIIE/ $\Delta I E \Theta N H \Sigma / \Sigma T O X O \Sigma / 0311 / 06$.

6 References

- [1] J. R. Black. *Electromigration: A Brief Survey and Some Recent Results*. In IEEE Transactions on Electron Devices, Vol. 16, No. 4, pp. 338-347, April 1969.
- [2] Y. Hoskote et al. *Failure Mechanisms and Models for Semiconductor Devices*. Report JEP122G, JEDEC Solid State Technology Association, Oct. 2011.
- [3] L. Huang et al. *AgeSim: A Simulation Framework for Evaluating the Lifetime Reliability of Processor-Based SoCs*. Proc. Design, Automation & Test in Europe Conference & Exhibition, pp. 51-56, March 2010.
- [4] H. Kim et al. *Use It or Lose It: Wear-Out and Lifetime in Future Chip Multiprocessors*. Proc. 46th IEEE/ACM Int'l Symposium on Microarchitecture, pp. 136-147, Dec. 2013.
- [5] M. Radetzki et al. *Methods for Fault Tolerance in Networks-on-Chip*. ACM Computing Surveys, Vol. 46, No. 1, Article No. 8, Oct. 2013.
- [6] M. Ramakrishna et al. *GCA: Global Congestion Awareness for Load Balance in Networks-on-Chip*. Proc. 7th IEEE/ACM Int'l Symposium on Networks-on-Chip, pp. 1-8, April 2013.
- [7] S. R. Sarangi et al. *VARIUS: A Model of Process Variation and Resulting Timing Errors for Microarchitects*. IEEE Trans. on Semiconductor Manufacturing, Vol. 21, No. 1, pp. 3-13, 2008.