

## ABSTRACT

Today, many-core and ultra-performing parallel architectures like Chip-Multiprocessors (CMPs) and Multi-Processor Systems-on-chips (MPSoCs) utilize Networks-on-Chips (NoCs) as their inter-tile communication infrastructure. NoCs are the preferable communication medium since they are able to overcome scalability and performance limitations that are common shortcomings in point-to-point connections, such as dedicated wires, and bus-based communication systems. Though the miniaturization of transistors has made the design and construction of CMP and MPSoC systems feasible, this technology scaling has come at the cost of increased vulnerability to wear-out, compromising the operational reliability of these systems.

Physical effects such as Electro-Migration (EM) and negative bias temperature instability, that are becoming more common due to transistor downsizing, may give rise to earlier transistor aging, increased electrical noise, elevated operational temperatures, and consequently eventual digital component breakdown. Communication links in NoCs are especially susceptible to faults due to the effects of EM. A single broken link can render the entire NoC as non-operational, as a routing algorithm oblivious to the presence of faulty links may not deliver messages to their destinations causing the NoC to stall completely. To overcome this detrimental outcome, NoC architects must design mechanisms to overcome the presence of such faulty network components, i.e. links. One major solution is to design appropriate fault-tolerant routing algorithms that can bypass faulty links in the NoC altogether, which will also be able to sustain relatively high throughput levels even with the presence of faulty links.

In this Thesis we propose HERMES, a fault-tolerant and load-balancing routing algorithm suitable for two-dimensional mesh-based NoC topologies. HERMES guarantees packetized message delivery in non-healthy NoCs which operate under a disconnected topological environment, while sustaining high-performance levels through graceful performance degradation in the presence of increased faulty link numbers. HERMES is a hybrid fault-tolerant routing algorithm: it utilizes deterministic routing such as dimension-order routing or 01TURN routing when faulty links are not present in a message's path, aiming to sustain high-performance, while it provides escape path selection in the vicinity of faults based on

up\*/down\* routing to deliver packets to their destinations in a deadlock-free mode, hence guaranteeing high NoC reliability.

HERMES was simulated under uniform random and transpose synthetic traffic patterns, with a range of virtual channel per port counts using wormhole flow-control, in order to determine its performance and behavior, utilizing two spatial faulty link placement scenarios: (1) random, and (2) hotspot faulty link distributions. When compared against ARIADNE, an existing state-of-the-art fault-tolerant routing algorithm, HERMES demonstrated up to 228.57% and 225% improvement in throughput with a random faulty link placement, while it showed up to 311.76% and 194% increase in throughput with a hotspot faulty link placement, under uniform random and transpose traffic pattern usages, respectively. HERMES was also tested using the Netrace benchmark suite demonstrating up to 38.83% improvement in network packet delivery latency when compared to ARIADNE.

Furthermore, HERMES' fault-tolerant scheme also includes a sub-network detection mechanism. This allows the discovery of non-communicating sub-areas and the determination of sub-network boundaries in case numerous consecutively spatially-placed faulty links cause the network topology to disconnect into disjoint router sets. With this sub-network detection mechanism we are able to provide sufficient information to the operating system in managing a CMP or MPSoC, so that they can utilize partitioned network topologies and archive higher core utilization even with large numbers of faulty links being present in their inter-tile interconnects.