

Department of Electrical and Computer Engineering
University of Rochester, Rochester, NY
Ph.D. Public Defense

Friday, December 2, 2016
10:00 AM
Computer Studies Building, Room 426

Energy-Efficient NoC Router Design with Adaptive Fault-Tolerance

Abstract

Energy-efficiency and fault-tolerance have become key challenges in the design of network on chip (NoC) using nanoscale technologies. As the communication backbone of chip multiprocessor (CMP), NoC consumes a significant portion of on-chip power which has already been one of the most critical constraints for chip design. Improving energy-efficiency of NoC thus becomes an imperative issue. On the other hand, due to the aggressive transistor scaling, on-chip defects increase at each new technology node. Unfortunately, NoC is susceptible to faults as NoC nodes highly depend on each other to establish a communication path. Failure of even one of these nodes can paralyze a large region of healthy cores that connected via it. Fault-tolerance thus becomes a great concern in NoC design. Current fault-tolerant NoC designs, however, either cannot provide enough protection from severe faults in the worst case or incur prohibitively high area, performance, and power overhead in the typical cases. This thesis aims to improve energy-efficiency and fault-tolerance of NoC at low cost by introducing novel designs of NoC routers.

This thesis starts with an eDRAM-based router buffer that improves router area and energy efficiency at the component level. As a key element that can greatly impact NoC performance, buffer accounts for a significant portion of router area and power. Instead of using traditional SRAM, we implement the router buffer with planar eDRAM to leverage its small size and low-power potential. We demonstrate that retention time of currently available eDRAM is much higher than the requirement for NoC. The implementation overhead is reduced by adopting a lightweight sense amplifier and a need-based refresh mechanism accordingly. Our design significantly reduces buffer area and power while maintaining similar performance as the SRAM-based buffer.

Then a NoC router with variable channel width is introduced to improve router energy efficiency at the architecture level. Based on the observation that short control messages account for a significant portion of NoC traffic, we use wide channel to transmit long data messages while dynamically split it into two narrower channels for short messages and shut down the unused channel to save energy. Experiment results show our proposed approach reduces NoC power consumption significantly under both real application traffic and synthetic traffic at the cost of negligible area overhead.

To bridge the competing goals of energy efficiency and fault tolerance, we propose an energy-efficient NoC router that exhibits strong fault-tolerance by leveraging channel slicing. This router has three identical router slices connected via internal sharing paths. When faults occur in any of them, resource sharing is enabled to enhance fault-tolerance. Channel slicing reduces the overhead of applying power gating, improving energy efficiency. These router slices can also be harnessed for on-demand TMR, to improve fault-tolerance further and reduce the need for deploying costly on-chip testers. With these techniques, our proposed router can tolerate much more faults than the state-of-the-art fault-tolerant NoC routers and consumes significantly less energy.