The Design and Use of High-Speed Transmission Line Links for Global On-Chip Communication

by

Aaron Carpenter

Submitted in Partial Fulfillment

of the

Requirements for the Degree

Doctor of Philosophy

Supervised by

Professor Michael Huang

Department of Electrical and Computer Engineering
Arts, Sciences and Engineering
Edmund A. Hajim School of Engineering and Applied Sciences

University of Rochester
Rochester, New York

2012

Abstract

As transistors approach the limits of traditional scaling, computer architects can no longer rely on the increase in density and core frequency to improve the overall system speed. Additionally, attempts to improve performance often result in disproportionately increased power and energy consumption. However, the increased performance and maximum frequency of the transistor allows us to build high-speed circuits specifically for on-chip communication. By incorporating the improving and emerging high-speed circuit technologies into the microprocessor design, it is possible to decrease the power and energy consumption, while simultaneously increasing system performance.

This thesis focuses on exploiting and analyzing the architectural opportunities provided by incorporating high-speed communication circuits, specifically on-chip transmission lines and simple high-speed transceivers. In broad terms, the transmission lines are used for a globally shared-medium on-chip interconnect, providing a low-latency, low-energy, packet-relay-free point-to-point link. Even a simple interconnect design can provide more than sufficient performance for small- to medium-scale chip multiprocessors. Additionally, with simple optimzi-ations exploiting benefits of a TLL shared-medium bus, it is possible to mitigate scalability limitations, and provide performance and energy benefits for larger-scale systems. For example, an atomic, low-latency bus provides opportunities to change the cache coherence substrate and optimize boolean data communication. This thesis will present and evaluate a number of these optimziations, and provide a final recommended design, showing performance and energy benefits with larger scale systems.