David H. Albonesi
School of Electrical and Computer Engineering
Cornell University
Ithaca, NY 14853
(607) 254-5473
albonesi@csl.cornell.edu

Philippe M. Fauchet and Eby G. Friedman
Dept of Electrical and Computer Engineering
University of Rochester
Rochester, NY 14627-0231
(585) 275-1487, (585) 275-1022
{fauchet,friedman}@ece.rochester.edu

Based on International Technology Roadmap for Semiconductors (ITRS) predictions, we develop a comprehensive comparison between optical and electrical interconnects for the design criteria of delay uncertainty, latency, power dissipation, and bandwidth density. We analyze an optical interconnect system consisting of three primary monolithically integrated on-chip components (light modulators, waveguides, and photodetectors) and assume an off-chip light source. Physical models have been developed to predict the operation of the optical components and to characterize their performance. The electrical interconnect is copper-based, and the wire size and inserted repeaters are optimized for the minimum delay, and minimum delay-power product, at different technology nodes from the ITRS Roadmap. From the comparative analysis of optical and electrical interconnect systems, optics exhibits a significant advantage in terms of latency, power dissipation and delay uncertainty at interconnect lengths as short as 2 to 7 mm. With the scaling of CMOS technology, the critical lengths for those metrics are well below the expected chip edge length. A key finding is that in highly scaled CMOS technologies, optical interconnects exhibit considerably reduced delay uncertainty compared to repeatered electrical wires, permitting the interconnect channel to be switched at a higher rate. However, the bandwidth density of single-wavelength optical interconnects is significantly lower than that of electrical interconnects due to the large dimensions of optical waveguides. Aggressive wavelength division multiplexing (WDM) of up to nine channels is necessary for optical interconnect to achieve a higher bandwidth density than electrical interconnect.

Based on this comparative analysis, we derive the requirements for silicon-based intra-chip optical interconnects and identify those areas where progress is urgently needed. These requirements can serve as guidelines in the development of optical devices. For example, the requirements suggest a focus on the development of high performance CMOS-compatible integrated WDM components.

The introduction of optical interconnects is also of high interest for computer architects. Based on the properties of optical interconnects developed from this investigation (most notably significant latency and delay uncertainty advantage in global interconnects), novel computer architectures are possible that fully exploit the advantages of optics. We have designed a new floorplan for high performance microprocessors consisting of multiple computation units and local caches. In a conventional floorplan with electrical interconnects, the hot computational units are often tightly packed to reduce inter-processor communication, creating excessively high on-die temperatures. Optical interconnects are used to manage the on-chip thermal budget. By using optics, computational cores can be distributed over the entire chip surface and surrounded by cooler cache memories, while maintaining cycle-level performance. With the improved thermal budget, an average speedup of over 50% is demonstrated across different benchmarks.

We further investigate other means by which silicon photonics can favorably disrupt the design of multi-core microprocessors. In particular, we find that the performance advantages of optical intra-chip technology can potentially simplify the cache coherency mechanisms that are often the most difficult part of the design to verify. The use of silicon photonics may permit designers to return to snoopy-based coherence protocols whereas electrical interconnects will require much more complex distributed directory schemes. Thus, successfully climbing the learning curve of integrating photonic elements, and putting in place the necessary library and tools infrastructures, significantly reduces design risk at the microarchitecture level.

Thus, topics ranging from new optical device models, rigorous electrical/optical circuit characteristics, and novel computer architectures will be discussed in an integrated way, providing enhanced insight into future trends in on-chip optical interconnect.