Optically-Connected Memory Systems for High-Performance Computing

Daniel Brunina, Caroline P. Lai, Ajay S. Garg, and Keren Bergman
Dept. of Electrical Engineering, Columbia University, New York, NY, 10027; daniel@ee.columbia.edu

Abstract: *We experimentally demonstrate a scalable memory architecture for high-performance computing using an optical interconnection network. Error-free communication is achieved between SDRAM and an emulated microprocessor using wavelength-striped 4×2.5-Gb/s streams across a 2×2 optical test-bed.*

Introduction

The limitations of main memory accesses have become a bottleneck for the scaling of high-performance computing systems. As such, the electronic interconnect between a processor and memory is a key consideration in overall system design [1]. Typical memory systems are based on synchronous dynamic random access memory (SDRAM), which is limited to clock frequencies an order of magnitude slower than those of high-performance processors. Therefore, a large number of SDRAM devices must be accessed in parallel to meet the high memory bandwidth requirements. This parallel data is transmitted across system boards over a double-pumped, path-length-matched, wide electronic bus.

Total memory capacity must scale with processor performance, further necessitating the need for additional SDRAM devices. However, to mitigate the inherent memory access latency, systems must minimize the physical distance between processors and main memory. Overall, the increasing number of SDRAM devices, higher bus data rates, and strict special design constraints are pushing the limits of electronic interconnects and threatening the performance gains of future computing systems [2].

An optically-connected memory architecture (Fig. 1), enables continued performance scaling due to its high-bandwidth capacity, energy-efficient bit-rate transparency, and time-of-flight latency. SDRAM parallelism and total capacity can scale to match future high-performance computing requirements without sacrificing energy-movement efficiency [3].

Experimental Setup and Results

The experimental setup presented here consists of two identical circuit boards that leverage an all-optical, wavelength-striped communication link through an implemented 2×2 all-optical network test-bed [4]. One circuit board acts as the microprocessor and the other as main memory. Each circuit board contains a field-programmable gate array (FPGA), four SDRAM modules, and transceivers capable of up to 4×2.5-Gb/s bidirectional communication. The transceivers on one board interface with optical modulators, creating a wavelength-striped 4×2.5-Gb/s optical stream. The stream propagates through the test-bed, is received using p-i-n-TIA receivers, and is transmitted to the second circuit board. This enables a communication link where all communication between the microprocessor and memory is performed all-optically.

Fig. 1: Block diagram of optically-connected memory.

The microprocessor implements a series of tests to validate complete memory functionality over the full memory address space, as well as overall network stability and reliability. Each test involves first writing to all memory addresses with a different bit pattern, and then reading back from all locations for verification. The read data is compared against expected bit patterns as it returns from memory over the optical test-bed.

Results and Conclusions

We experimentally demonstrate an emulated microprocessor communicating with SDRAM modules over an implemented optical network test-bed. Error-free operation is achieved by successfully writing and reading one gigabit of data without bit mismatches (effective memory bit-error rate less than 10^{-9}).

References

[1] R. Yavatkar, HPCA-12 (2006).
[2] W. Wulf et al., ACM SIGARCH **23** (1) (1995).
[3] Y. Katayama et al., HPCA 46-50 (2007).
[4] A. Shacham et al., PTL **17** (12) (2005).