

# Experimental Demonstration of Optically-Connected SDRAM

**Daniel Brunina, Ajay S. Garg, Howard Wang, Caroline P. Lai, and Keren Bergman**

*Department of Electrical Engineering, Columbia University, 500 W. 120<sup>th</sup> St., New York, New York 10027  
daniel@ee.columbia.edu*

**Abstract:** A four-channel, 2.5-Gb/s, all-optical WDM link is established between SDRAM and an emulated CPU. Data integrity and error-free performance are verified with a sequence of SDRAM write and read operations.

©2010 Optical Society of America

OCIS codes: (200.4650) Optical interconnects, (060.4510) Optical communications, (200.0200) Optics in Computing;

## 1. Introduction

The recent trend towards the multiplication of on-chip processing cores in microprocessors and of the number of chips per board is resulting in the on-chip communications infrastructure becoming a key consideration in the design of high-performance computing systems [1]. Current electronic interconnects provide a critical communication path between main memory and processor; however, the need to drive signals at increasingly high speeds has resulted in a dramatic increase in power dissipation due to the high-frequency attenuation characteristics of electronic wires. Combined with the limited available interconnect density, the increasing power demands are so acute that this requirement is a crucial limitation to the continued scaling of computing systems [2], [3]. Ultimately, the performance of electrical memory connections will be limited by metal wire length, skew, clock frequency, and pin count.

Fig. 1 shows the physical wiring of a single memory channel bus between four dual in-line memory modules (DIMMs) and the memory controller of an Intel chipset featuring a typical electronic interconnect. Each DIMM may have up to 240 pins, and the path lengths of both data and control signals must be matched. In the current memory access design paradigm, the DIMMs must be placed as close to the memory controller as possible in order to minimize the power dissipation and latency associated with driving and buffering long electrical wires. Minimizing the wire and buffer delays is a critical design constraint due to the fact that a typical processor clock operates on the order of several gigahertz while a memory clock operates on the order of only hundreds of megahertz. This speed disparity has exacerbated in recent years, as the number of processor cycles per memory access has steadily increased, resulting in a memory wall that essentially negates any performance improvements gained by increasing processor speeds [4].

In this work, we experimentally demonstrate a photonic memory access system that exploits the parallelism of wavelength-division multiplexing (WDM). Contrary to traditional electronic designs, the large bandwidth-distance product gained by an optical interface to memory does not require the DIMMs to be as close to the processor as physical limitations permit. Rather, leveraging optics allows the memory modules to be placed on a separate board to drastically reduce processor-to-memory electronic wiring, as well as permit other critical hardware to be located close to the processor [5]. The separate memory board can also provide more area for a greater number of DIMMs, which can increase both memory capacity and bandwidth via data parallelism without sacrificing the power efficiency of the memory interface.

For the purpose of this demonstration, the DIMM is implemented on a development board containing a field-programmable gate array (FPGA) connected to two independent banks of 64Mx32 DDR2 SDRAM totaling 512 MB (four chips of MT47H64M16BT-37E). A separate FPGA-based device models a CPU with an on-chip memory controller. The function of the CPU is to generate read and write memory accesses while continually checking for errors. The two FPGAs communicate through a 4x2.5-Gb/s optical link that is transparent to the CPU, memory controller, and main memory.

## 2. Electronic Memory Access Architecture

CPU-DIMM communication is handled by a device called the memory controller. A majority of personal computers (PCs) and servers utilize memory controllers that communicate with DIMMs consisting of synchronous dynamic random access memory (SDRAM). Traditionally, these memory controllers are in chips separate from the CPU chip, such as most Intel chipsets (Fig. 2). Off-chip memory controllers reduce the size and cost of the CPU at the expense

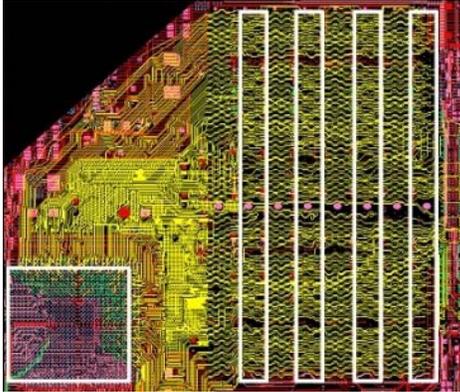


Fig. 1. Electrical wiring of a single memory channel bus between a DIMM and a memory controller of an Intel chipset. White rectangles correspond to DIMM locations.

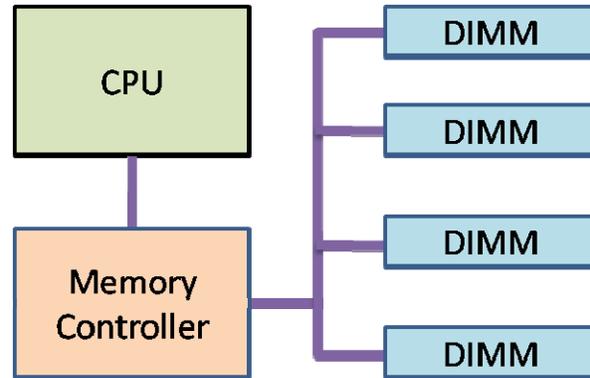


Fig. 2. Logical architecture of traditional CPU, memory controller, and DIMMs. Here the memory controller is located off-chip with respect to the CPU.

of limited performance due to electronic interconnect limitations. Recently, the trend is to increase overall CPU-memory controller communication bandwidth by placing the memory controller on-chip [6].

In order to access main memory, the CPU generates read or write requests and transmits them to the memory controller. The memory controller translates each request into a sequence of SDRAM-specific row and column access commands, which ultimately activates specific locations within the multiple arrays of storage elements on an SDRAM module. Since the SDRAM clock is relatively slow as compared to the CPU clock, DIMMs are designed to increase parallelism by striping data through the SDRAM banks and accessing many data words simultaneously. These can be accessed with shared SDRAM commands, therefore minimizing the number of commands necessary to load or store each unit of data. Modern SDRAM designs exploit this parallelism with a technique called double data rate (DDR), or double pumping. A DDR DIMM fetches and buffers multiple consecutive words of data simultaneously without extra SDRAM commands. The memory I/O bus transmits data twice per clock cycle -- once at the rising edge and once at the falling edge -- and is also clocked at a frequency multiple times faster than the internal memory clock. The latest commercially available DDR technology, DDR3, acts on eight words simultaneously and clocks the I/O bus at four times the memory clock.

Optical interconnects map well to this trend of increased bandwidth through data parallelism and high-speed I/O busses. An optical fiber does not suffer the same performance or energy penalties as electronic interconnects, which allows DDR technology to continue scaling the I/O frequency. Furthermore, by exploiting the parallelism of WDM, several SDRAM chips or entire DIMMs can be accessed in parallel over a single fiber to increase the memory bandwidth. Overall, replacing the wide electronic memory bus with an optical fiber can not only increase bandwidth, reduce system wiring complexity, and lower energy consumption, but also allow the current trend of DDR technology to continue scaling into the future.

### 3. Experimental Demonstration and Results

In order to demonstrate the feasibility of an optical interface to memory, the following experiment emulates a CPU-Memory system with two identical high speed Altera FPGAs. Each FPGA is on a separate board, as seen in Fig. 3, with focus placed on implementing the hardware necessary for main memory transactions: CPU, memory controller, I/O bus, and SDRAM. The CPU and memory controller are synthesized together on one FPGA, creating an on-chip memory controller, along with an optical format module (OFM) to convert between the large controller-SDRAM bus and a four-channel 2.5-Gb/s WDM optical link. The OFM is also synthesized onto the second FPGA, which is electronically connected to two independent banks of DDR2 SDRAM. The CPU, memory controller, and SDRAM are unaware that the optical link transparently replaces the electronic I/O bus.

The memory controller is configured as a single 32-bit channel that utilizes a 500-MHz memory bus. With two data transfers per clock cycle, the peak memory bandwidth is 32 Gb/s. The bus from the memory controller has been modified to packetize the SDRAM commands and read/write data for transmission over the optical link, which provides bandwidth equalization between the memory controller and the four 2.5-Gb/s FPGA transceivers. A traditional memory controller is designed to interface with a wide electronic bus that is directly wired to every SDRAM module. The data packetization requires the memory driver hardware to be moved from the memory controller onto the SDRAM board, due to the requirement that relative timing of the memory controller signals is to

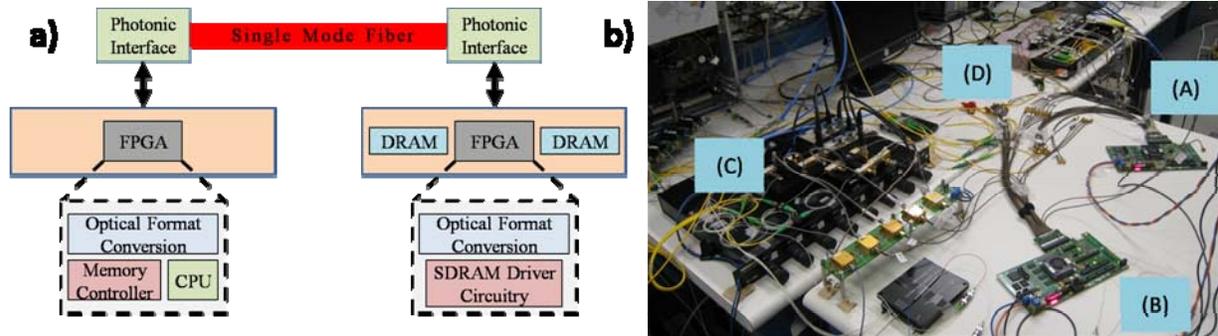


Fig. 3. a) Block diagram of optically-connected memory experiment. CPU, memory controller, and optical format conversion modules are synthesized onto an FPGA (left). A second optical format conversion module is synthesized onto a second FPGA (right), which is electrically connected to two independent SDRAM banks. b) Experimental setup, showing FPGAs (labeled A and B), optical modulation equipment (C), and optical receivers (D).

be maintained. However, the hardware for this is minimal and does not introduce significant complexity to the design of a remote optically-connected memory board.

To experimentally validate the optically-connected memory system, the CPU is emulated as a module that generates a pattern of memory write and read requests and continually verifies the received read data to detect errors. The stored data in the first phase of testing is a binary pattern of all ones. In the second phase, the CPU generates a pseudo-random sequence of data based on the destination write address. These tests have the effect of verifying full memory functionality, including access to full memory address space with write and read capabilities, as well as demonstrating the stability of the optical memory system by transporting over a gigabit of data in both directions across the optical link. The result is an effective memory bit error rate (EMBER) of less than  $10^{-9}$ .

#### 4. Conclusion

In this work, we experimentally demonstrate the functionality and viability of optically-connected SDRAM. A CPU with on-chip memory controller is emulated on a high-speed FPGA, which communicates with commercial SDRAM modules over four 2.5Gb/s optical channels. Significant opportunities exist for improving the optical memory system, including the use of optical switching for implementing multiple optically-connected memory nodes as well as integrating SDRAM modules and optical components onto a single board. This work suggests that future systems may leverage optics to locate main memory remotely from the processor without reducing bandwidth or sacrificing energy, thus improving overall memory access designs.

#### 5. References

- [1] S. Vangal et al., "An 80-Tile 1.28TFLOPS Network-on-Chip in 65nm CMOS," ISSCC Dig. Tech. Papers, pp. 98-99, Feb. 2007.
- [2] R. Ho, K. W. Mai, M. A. Horowitz, "The future of wires," *Proceedings of the IEEE*, vol. 89, no.4, pp. 490-504, Apr 2001.
- [3] J. Shalf, "The new landscape of parallel computer architecture," *Journal of Physics: Conference Series*, vol. 78, no. 1, July 2007.
- [4] W. A. Wulf, S. A. McKee, "Hitting the memory wall: implications of the obvious," *ACM SIGARCH Computer Architecture News*, vol. 23, no. 1, pp. 20-24, March 1995.
- [5] Y. Katayama, A. Okazaki, "Optical interconnect opportunities for future server memory systems," *Proceedings of the 2007 IEEE 13th International Symposium on High Performance Computer Architecture*, pp. 46-50, February 2007.
- [6] B. Stackhouse, B. Cherkauer, M. Gowan, P. Gronowski, C. Lyles, "A 65nm 2-Billion-Transistor Quad-Core Itanium<sup>®</sup> Processor," *Solid-State Circuits Conference, 2008. ISSCC 2008. Digest of Technical Papers. IEEE International*, pp.92-598, 3-7 Feb. 2008.