

Design and Evaluation of an Optical CPU-DRAM Interconnect

Amit Hadke ^{*}, Tony Benavides [†], Rajeevan Amirtharajah [†], Matthew Farrens ^{*}, Venkatesh Akella [†]

[†] Department of Electrical & Computer Engineering

^{*} Department of Computer Science

University of California, Davis, CA - 95616

Email: akella@ucdavis.edu

Abstract—We present OCDIMM (Optically Connected DIMM), a CPU-DRAM interface that uses multiwavelength optical interconnects. We show that OCDIMM is more *scalable* and offers *higher* bandwidth and *lower* latency than FBDIMM (Fully-Buffered DIMM), a state-of-the-art electrical alternative. Though OCDIMM is more power efficient than FBDIMM, we show that ultimately the total power consumption in the memory subsystem is a key impediment to scalability and thus to achieving truly *balanced* computing systems in the terascale era.

I. INTRODUCTION

In 1967 Gene Amdahl defined a balanced system as one that provided equal numbers of bytes of memory, bits/second of I/O capacity, and instructions per second [1]. More recently, a blue-ribbon NSF panel chaired by Atkins [2] updated the definition to be one byte of memory per FLOP, and a memory bandwidth greater than or equal to one byte/s/FLOP/s. Thus, building a balanced multicore processor will require a significant increase in both available memory and bandwidth.

The standard technique for increasing memory capacity is to use denser DRAM chips. However, due to electrical signaling constraints, the maximum number of supportable chips per channel has been falling over successive generations of DRAMs. In addition, the traditional approaches to increasing memory bandwidth (using higher clock frequencies and/or a wider parallel bus) do not scale well because of the limitations on the number of pins on a package and the inherent complexity of running wide parallel buses at very high clock frequencies.

The FBDIMM (fully-buffered DIMM) architecture [3] is a recent commercially available approach to providing more bandwidth and increased memory capacity. It accomplishes this by using a narrow high-speed *point-to-point* interface between the memory controller and the memory modules (DIMMs). Unfortunately, because of the limitations of the store-and-forward protocol employed within the FBDIMM, as more DIMMs per channel are added the latency increases substantially. Thus, although FBDIMM provides a significant improvement over DDR2/DDR3 based interfaces, it does not scale well beyond 8 DIMMs per channel.¹ FBDIMM by

¹There is a variable latency mode of operation supported by FBDIMM to help deal with this problem, but it makes the memory controller significantly more complex and in most cases the fixed-latency mode of operation is preferred.

itself will not be able to provide both the increased memory capacity necessary to keep up with the increasing working sets of future applications and the bandwidth scalability in order to maintain a balanced system that will be required in future multicore chips.

Optical communication links are known to offer significantly higher bandwidth over longer distances, and during the past few decades they have made rapid inroads into applications such as local area networks, storage networks and rack to rack and board to board communication links. However, cost, optical losses, and incompatibility with nanoscale CMOS has kept them away from inter-chip and intra-chip applications. Fortunately, recent developments in CMOS-compatible nanoscale silicon photonic integrated circuits [4], [5] make optical interconnects a viable alternative in future high-performance computing systems. Researchers at MIT [6] have demonstrated the feasibility of monolithic silicon photonics technology suitable for integration with standard bulk CMOS processes, which reduces costs and improves opto-electrical coupling.

In [7], we proposed an extension to the FBDIMM architecture called OCDIMM (Optically Connected DIMM) that replaces the electrical links between the memory controller and the DRAM modules with a wavelength division multiplexed (WDM-based) optical bus. In this paper, we extend the analysis reported in [7] in two ways. First, we evaluate the impact of using DDR3 memory modules, and second we investigate the power consumption of OCDIMM based memory systems. We show that on average an OCDIMM based system consumes about 15% to 20% less power than a corresponding FBDIMM based system. However, even with optical interconnects, we estimate the total power consumption of a 64GB OCDIMM based memory system to be around 182W (for a DDR2-667 based system) and 142W (for a DDR3-1333 based system), which is quite significant and presents a serious challenge to building *balanced* computing systems in the terascale era. We investigate the sources of the power consumption which in turn indicate potential avenues for reducing the power consumption in an optically connected memory system.

The rest of this paper is organized as follows. First, we present an overview of the OCDIMM architecture. This is followed by a description of our experimental methodology, including how OCDIMM was modeled in the DRAM simulator. We evaluate the potential of OCDIMM by comparing

the latency, bandwidth, scalability and power consumption of OCDIMM and FBDIMM for both DDR2 and DDR3 based DIMMs, and finally we conclude the paper with an overview of related work and directions for future work.

II. OVERVIEW OF OCDIMM

The reader is referred to the excellent textbook [8] and tutorial paper [9] for more details on the FBDIMM architecture, upon which OCDIMM is based.

A. FBDIMM and OCDIMM

The FBDIMM memory architecture replaces the shared parallel interface between the memory controller and the DRAM chips with a point-to-point serial interface between the memory controller and an intermediate buffer, the Advanced Memory Buffer (AMB). The on-DIMM interface between the AMB and the DRAM chips is identical to that seen in DDR2 and DDR3 systems, and will not be discussed further.

The serial interface is split into two unidirectional buses, one for read traffic (the northbound link) and one for command/write traffic (the southbound link). FBDIMM uses a packet-based protocol that bundles commands and data into frames that are transmitted on the channel and then converted to the DDR2/DDR3 protocol by the AMB. Since DIMMs are connected in point-to-point fashion, a channel becomes a multi-hop store and forward network. All frames are processed by the AMB to determine whether the data and commands are addressed to a local DIMM. This is a serious limitation in FBDIMM, and was a prime motivating factor in developing OCDIMM.




Fig. 1. Top Level System

Figure 1 shows a high-level diagram of a system using OCDIMM. The key difference between this and FBDIMM is that the northbound and southbound buses are replaced

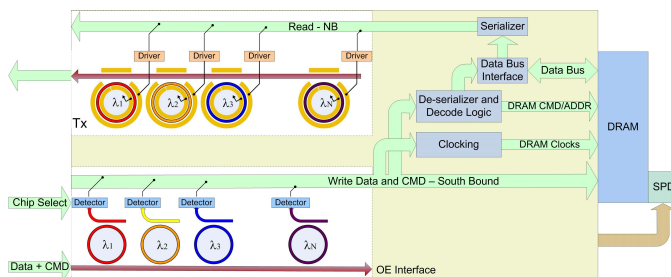


Fig. 2. OMC Block Diagram

by optical fibers. To support this the AMB on each DIMM is replaced with an Optical Memory Controller (OMC), which handles all communication between the DRAM and the bus. The details of the OMC are shown in Figure 2. We assume electrical to optical (E/O) converters on the integrated memory controller, which convert the electrical signals to optical signals that are transported over the optical waveguides and converted back to the electrical domain at the OMC. We also assume that each waveguide can transport multiple wavelengths (up to a maximum of 64), although we will show in the results section that it is not necessary to have 64 wavelengths to reap the benefits of an optical interconnect. In essence we have created a single-hop (broadcast) bus instead of a multi-hop store and forward network, which significantly decreases the latency.

B. Memory transactions in OCDIMM

Read and write data is organized using a packetized frame relay protocol which is realized using multiple transmissions on the optical bus (the number of transmissions can be reduced if there are more wavelengths available). Writing data to the DRAM is simple - depending on the DRAM mode, a RAS command followed by some number of CAS commands are sent to the target DIMM, interspersed with the data on the write channel. Once the commands and data are converted to the electrical domain via the E/O/E interface in the target domain, the operation is exactly the same as with a standard DDR2/DDR3 device and FBDIMM. No modification is necessary to the existing DRAM devices.

Reading data from a given DIMM is slightly more complicated, because the read subchannel is shared. Once a DIMM is ready to send data, it has to acquire the read subchannel since it might be in use by another DIMM. In general, this requires an arbiter, but to keep the design simple we initially assume that the memory controller statically schedules the read transactions such that there are no bus conflicts.

OCDIMM uses the same frame length for the northbound and southbound buses as that used by FBDIMM. Because we are now working with a 10 Gbps optical data rate, this will increase the OCDIMM controller clock to DRAM clock ratio to 15:1 (up from the 6:1 ratio in FBDIMM). This increased ratio allows the transfer of 32 bytes per frame on the northbound lane while also transferring six commands or one command and 16 bytes of data per write on the southbound lane.

III. EXPERIMENTAL METHODOLOGY

We modified the DRAM simulator (DRAMsim) from the University of Maryland [10] to model the OCDIMM. We used the DDR2-667 and DDR3-1333 to model latencies within the DIMM based on datasheets from Micron. All DIMMs are symmetrical and have eight banks with an eight byte channel width. We used ranks to represent DIMMS, as done in [9]. Each DIMM has only one rank, and to add a DIMM we add a rank in the simulator configuration.

The FBDIMM memory controller is set to run at 6x DDR2-667Mbps (i.e at 4002 MHz). Each AMB uses

up/down buffers for write/read requests, and the buffer count is set to the number of banks on a DIMM (eight). The flight time from the controller to the first DIMM is assumed to be zero, while the time from DIMM-to-DIMM is modeled as 1.5ns (this includes the minimum resample mode delay). Thus all latencies and bandwidths are measured at the point right before the first DIMM. For simulations of OCDIMM, the DIMM-to-DIMM flight time is set to 300ps, which includes the transmission delay for a distance of up to 2cm, O/E/O conversions, and skew adjustment. All optical buses operate at 10GHz.

OCDIMM is modeled in two different ways:

Frame cycle = DRAM cycle : The time required to send a frame is equal to one DRAM cycle. (For DDR2-667Mbps, frame cycle = DRAM cycle = 3ns). Depending upon the wavelengths (λ) the amount of data that can be transferred in a frame varies. For example, when $\lambda = 28$ for northbound and 20 for southbound, the northbound DATA frame carries 32 bytes, while the southbound COMMAND frame carries six commands and the southbound COMMAND+DATA frame carries a single command and 16 bytes of data. Similarly, when $\lambda = 16$ for both northbound and southbound channels, the northbound DATA frame carries 16 bytes, the southbound COMMAND frame carries three commands, and southbound COMMAND+DATA frame carries a single command and eight bytes of data.

Frame cycle \neq DRAM cycle : The frame time is varied according to λ . To model this correctly, when the frame cycle is less than the DRAM cycle extra delays are added in the AMB to ensure that it issues commands on the correct DRAM clock cycle.

Every transaction is assigned a uniformly distributed random address. We clubbed rank address bits with channel address bits, so as to map consecutive cachelines to different DIMMs. This gives maximum possible parallelism and deeper channels can schedule more transactions at any given time. We also used a random, uniformly distributed request arrival rate. This workload attempts to schedule as many transactions as possible, hence it is used in the experiments to measure maximum bandwidth with $\approx 100\%$ utilization. Each transaction reads or writes a cacheline of 64 bytes. Read-to-write transaction ratio was kept at 2:1 as most of the workloads are observed to have more read request than writes (Page 550-551 of [8]).

IV. RESULTS AND DISCUSSION

Figure 3 demonstrates the key benefits in terms of the improvement in sustained bandwidth for OCDIMM. We plot the average latency versus the sustained bandwidth for different configurations of FBDIMM and OCDIMM with DDR2-667 DIMMs. The data was generated by increasing the traffic injection rate, and the latency was measured from the time when read/write transaction was entered into BIU (bus interface unit) queue to when it is marked "complete" (i.e., all commands corresponding to the transaction are complete.) We assume 8 DIMMs for the FBDIMM case, which corresponds to its maximum capacity. For FBDIMM,




Fig. 3. Comparison of Sustained Bandwidth of FBDIMM and OCDIMM Configurations

we show both the fixed and variable latency modes of operation. As expected, variable latency mode is better than the fixed-latency mode and the sustained bandwidth is about 5.33 GBytes/s.

Wavelength-division multiplexing allows the transmission of multiple bits of data simultaneously on the same physical fiber by modulating them on different wavelengths. Therefore, we evaluate four different configurations of OCDIMM. First we show a *baseline* configuration which uses 32 and 64 wavelengths respectively, partitioned equally between northbound and southbound channels. Next we explore an advanced configuration that allocates dedicated wavelengths to different DIMMs. The configuration *8C2D* corresponds to a configuration where a total of 16 DIMMs are partitioned into 8 sets, with each set of two DIMMs allocated 8 wavelengths. The configuration *8C8D* corresponds to a total of 64DIMMs, partitioned into 8 clusters with each cluster of 8 DIMMs sharing 8 wavelengths.

The graphs clearly indicate two key benefits of OCDIMM - first, it is scalable to 64 DIMMs, while FBDIMM is limited to just 8 DIMMs. This means one can get much higher memory capacity with OCDIMMs. Second, significantly higher bandwidth can be sustained by the OCDIMM configuration. Also, note that even a baseline OCDIMM configuration with 32 wavelengths outperforms the variable mode FBDIMM.

Next we evaluate the potential of OCDIMM to reduce memory latency. Figure 4 shows the latency distribution for FBDIMM configuration operating in the variable latency mode with 8 DIMMs, while Figure 5 shows the latency distribution for OCDIMM configuration with 16 DIMMs and a total of 32 wavelengths partitioned equally between northbound and southbound channels. We show the impact on both read and write latency in each case. Clearly, OCDIMM has a significant advantage, especially for READ operations (which are more critical.)

We summarize the key results in Table I. We assume 16

TABLE I
KEY PERFORMANCE RESULTS

	Average Latency (ns)	Sustained Bandwidth (GB/s)
FBDIMM (DDR3-1333, variable latency)	99	9.54
FBDIMM (DDR3-1333, fixed latency)	102	9.3
OCDIMM (32 wavelengths)	66	15
OCDIMM (64 wavelengths)	61	17

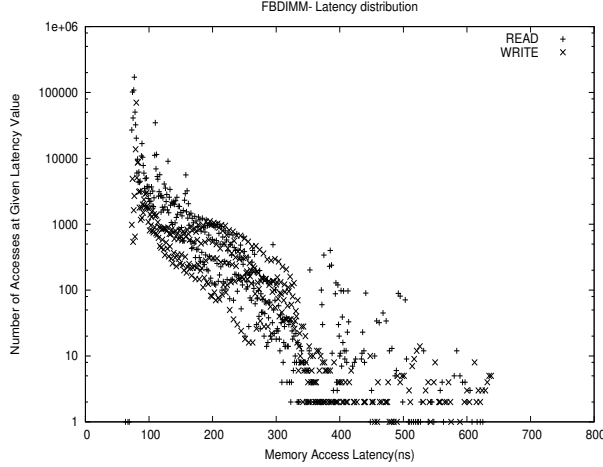


Fig. 4. Latency Distribution for FBDIMM with 8 DDR2-667 DIMMs and variable latency mode operation using gcc benchmark from SPEC CPU2000

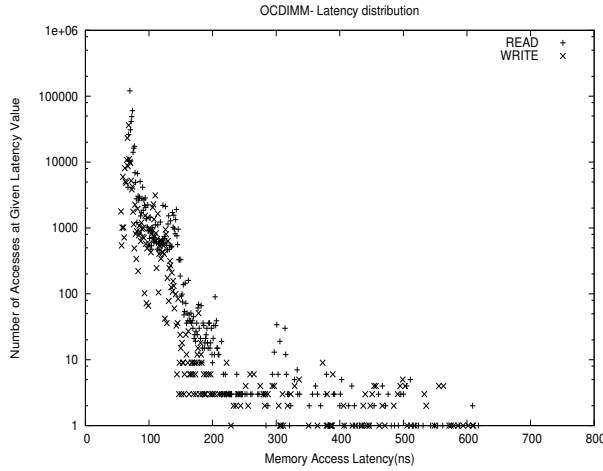


Fig. 5. Latency Distribution for OCDIMM with 16 DIMMs, DDR2-667 with 16 wavelengths for northbound and 16 wavelengths for southbound channel using gcc benchmark from SPEC CPU2000

DIMMs per channel for the OCDIMM configurations and 8 DIMMs per channel for the FBDIMM. For OCDIMM the wavelengths are assumed to split equally between northbound and southbound channels.

The results indicate that, when compared to FBDIMM, on average OCDIMM can improve the bandwidth by around 79% and reduce latency by 38%.

V. POWER ANALYSIS OF OCDIMM

We focus on the power consumption of the memory subsystem shown in Figure 1. We assume 65nm technology node and a 4GHz memory controller. The sources of power consumption include the DRAM chips (common for both the OCDIMM and FBDIMM), the AMB for the FBDIMM, the OMC for the OCDIMM, the coupling and fiber losses

and the power consumed by the optical modulators and detectors including the transceivers. As shown in Figure 2 the OMC components consist of the serializer, modulator, deserializer, detector and optical clocking source. We estimate the power consumption for the optical components used in the OCDIMM architecture based on the data for analogous components reported in [11] and [6].

We use Micron’s power calculator [12] to estimate the power consumed by the DDR2/3 DRAM chips based on the activity statistics obtained from the DRAMSim for a given memory trace. The AMB power consumption was obtained from DRAMSim’s FBDIMM model. We also model the situation when the DIMMs are *idle*. For FBDIMM, the *idle* time represents the time that the DRAM is in a known CKE power down state. This shuts down the output drivers and forces the memory controller to activate periodic refresh intervals approximately every 7.6 μ s. We assume that DRAM is in the *idle* state for 5% of the total system run time. When analyzing the *idle* time for the OMC we assume that the modulators, serializers, etc are also in a known powerdown state, thus reducing system power at IDLE time. There are some differences FBDIMM and OCDIMM when a given DIMM is idle. In the case of FBDIMM, the AMB is still passing data to other *active* DIMMs, even though the host DIMM is idle. For the OCDIMM case, when the DIMM is in an *idle* state, the OMC related to the DIMM does not have to relay data to its neighbour and hence can be placed in a power down state.

For DDR3-1333 we assume a low power DDR chip which runs at 1.35V. We assume that the onchip coupling loss $P_{couple,emtofiber}$ and $P_{couple,modimm}$ are each 1dB. The loss in silica fiber (waveguide) is 0.5–5 dB/cm. Focusing on using ring modulators instead of MZ modulators, the modulator insertion loss, $P_{modulator1}$, is 0.5dB, while the modulator driver power consumption, $P_{modulator2}$, is 70 fJ/b. For our external off-chip laser source, the Laser Power, P_{Laser} , is 300 fJ/b. The minimum detector power for the photodiode $P_{detector}$ is 12 fJ/b.

There are some additional components which will add to our total estimate of back-end electrical/photonic power analysis - these components include the serializer, deserializer, and optical clocking source, which add up to $P_{misc,lect} = 80$ fJ/b. An additional 100 fJ/b needs to be added to the equation to compensate for the extra thermal energy ($P_{thermal}$) that is used by the optical components. Thus the total input optical power, $P_{optical}$, is $N * P_{det} * 10^{(P_m + 2 * P_{couple}) / 10}$ [13] and

the total electrical power is $P_{elect} = N * (P_{md} + P_{misc,lect} + P_{laser} + P_{thermal})$, where N is the number of optical links (32 northbound and 32 southbound links when there are 64 wavelengths available).

After the energy unit conversion based on a 4GHz execution cycle rate, we get a total $P_{tot} = P_{elect} + P_{optical} = N * (1.509 \text{ mw/b})$. Based on this, we estimate the total additional power added by OMC on southbound channel to be 46 mw/DIMM and our northbound channel power estimate to be 64.4 mw/DIMM.

TABLE II
OMC POWER EQUATIONS PER DIMM

Power Component	Equation
P_{elect}	$P_{modulator} + P_{misc}$
P_{misc}	$P_{serializer} + P_{deserializer} + P_{other}$
$P_{optical}$	$P_{det} * 10^{((P_{modulator} + 2 * P_{couple} + 0.5 * 5) / 10)}$
P_{tot}	$P_{elect} + P_{optical} + P_{misc}$

Figure 6 and Figure 7 compare OCDIMM to FBDIMM with respect to power, for a DDR2 and a DDR3 based system, respectively. As expected, OCDIMM requires less power because the OMC does not use store and forward logic. If only 1 DIMM is populated, OCDIMM may have a high power consumption due to the additional overhead in the electrical-to-optical conversion and back - however, as the capacity increases, OCDIMM delivers substantial power savings (on average about 15% to 19%) compared to FBDIMM. One could argue that in absolute terms the power consumption of OCDIMM, especially at larger capacities (eg: 64GB), is quite high, but given the higher bandwidth we believe it is justifiable in high performance computing applications. In addition, opportunities for power optimizations can be seen by analyzing the sources of the power consumption as shown in Figures 6 and Figure 7.

It is interesting to note that the power consumption due to the optical overhead (modulators, laser source and detectors) is negligible compared to the power consumed by the OMC and the DRAM chips. Clearly the OMC has to be optimized in conjunction with the memory controller to reduce the power consumption. With the modulators and lasers taking up a good portion of the power envelope, an effort should be put forth that focuses on better modulation materials and external laser generation to reduce the overall system power requirements.

VI. RELATED RESEARCH

Recently, HP researchers proposed a 3D manycore architecture called Corona [14] that uses nanophotonic communication for both inter-core and off-stack communication to the I/O devices. In addition, MIT researchers [6] presented a new monolithic silicon technology suited for integration with standard bulk CMOS processes and have developed a processor memory network architecture for future manycore systems based on an optoelectrical global crossbar. Cornell researchers [11] have also described a methodology to lever-

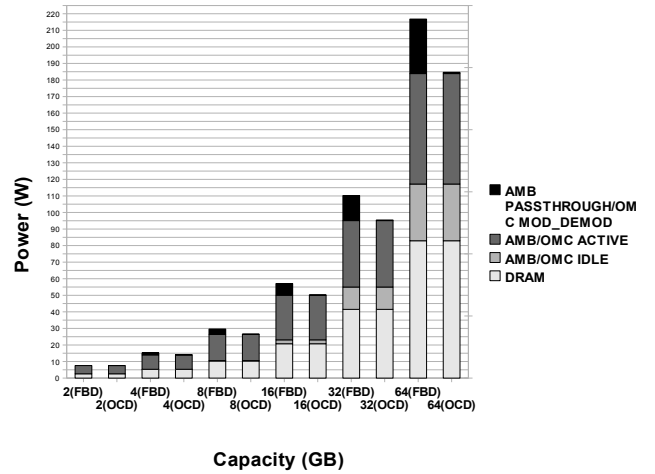


Fig. 6. Comparison of OCDIMM(OCD) and FBDIMM(FBD) Power Consumption for DDR2-667 Based DIMMs. Random traces with a fixed latency single channel open page mode with greedy scheduling and queue depth of 16. Capacity(number of DIMMS on channel) is varied on the X axis and the Y axis plots average power.

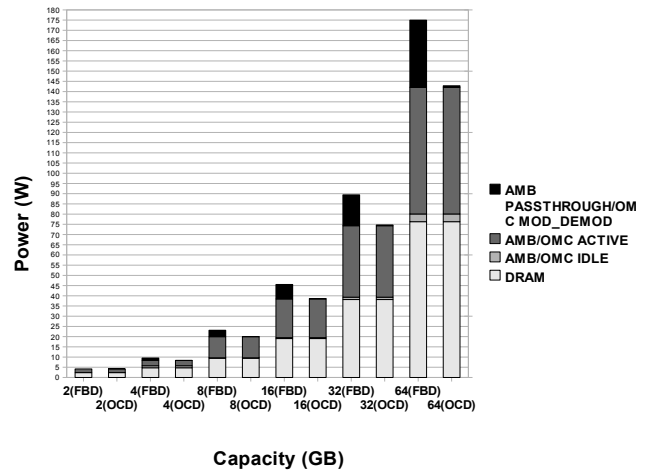


Fig. 7. Comparison of OCDIMM(OCD) and FBDIMM(FBD) Power Consumption for DDR3-1333 Based DIMMs. Random traces with a fixed latency single channel open page mode with greedy scheduling and queue depth of 16. Capacity(number of DIMMS on channel) is varied on the X axis and the Y axis plots average power.

age optical technology to reduce power and latency in on-chip networks suitable for bus-based multiprocessors.

Drost et. al. [15] point out the technological challenges to building a flat-bandwidth memory hierarchy in CMOS-based systems, and offer proximity communication as a way to provide both high bandwidth and high capacity simultaneously. We propose to address some of the same challenges using DWDM-based optical interconnect instead.

The Multi-Wavelength Assemblies for Ubiquitous Interconnects (MAUI) has been examining the idea of interfacing fiber to the processor (FTTP) and how the cost will affect the overall penetration of the market [16]. An older research project, the High Speed Opto-Electric Memory

System (HOLMS) interface, proposes an optical connection between the memory system and the processor, but acknowledges the need for a controller-to-DIMM interface to take advantage of the known optical networking success in the telecommunications industry [17].

Haas and Vogt introduced FBDIMM technology in [3], and researchers at the University of Maryland evaluated the potential of FBDIMM on real workloads in [9]. The controller-to-DIMM interface at the heart of FBDIMM, coupled with the recent advances in the design of nanoscale modulators and detectors with support for multiple wavelengths [4], [18]–[20] which can be used to realize the optical interconnect structures inherent in the OCDIMM architecture, led us to our design.

VII. CONCLUSIONS AND FUTURE WORK

We have proposed a simple extension to the FBDIMM architecture that takes advantage of multiwavelength optical interconnects. The resultant architecture, called OCDIMM, is found to improve capacity, reduce latency and increase bandwidth. As pointed out by HP researchers [4], DWDM based nanophotonics offers an additional dimension of concurrency that can be exploited in innovative ways to further improve the bandwidth and latency of an optical interconnect. We are investigating more advanced memory controller designs based on this idea. We find that power consumption in the OMC (the equivalent of the AMB in a FBDIMM) has to be reduced significantly to make OCDIMM truly scalable to large capacities, i.e. beyond 64GB. We are exploring the possibility of sharing an OMC by a set of DIMMs and redesigning the OMC from scratch instead of modifying the existing AMB.

VIII. ACKNOWLEDGMENTS

This project was supported in part by CITRIS Seed Funds Under Grant 3-34PICOB.

REFERENCES

- [1] G. Bell, J. Gray, and A. Szalay, "Petascale computational systems," *Computer*, vol. 39, no. 1, pp. 110–112, 2006.
- [2] D. Atkins, K. Droegmeier, S. I. Feldman, H. Garcia-Molina, P. Messina, and J. Ostriker, "Revolutionizing science and engineering through cyberinfrastructure - report of blue ribbon panel on cyberinfrastructure," National Science Foundation, Tech. Rep., 2003. [Online]. Available: <http://www.nsf.gov/od/oci/reports/atkins.pdf>
- [3] J. Haas and P. Vogt, "Fully-buffered dimm technology moves enterprise platforms to the next level," *Intel Technology Magazine*, 2005.
- [4] R. Beausoleil, P. Kuekes, G. Snider, S.-Y. Wang, and R. Williams, "Nanoelectronic and nanophotonic interconnect," *Proceedings of the IEEE*, vol. 96, no. 2, pp. 230–247, Feb. 2008.
- [5] M. Lipson, "High performance photonics on silicon," *Optical Fiber communication/National Fiber Optic Engineers Conference, 2008. OFC/NFOEC 2008. Conference on*, pp. 1–3, Feb. 2008.
- [6] C. Batten, A. Joshi, J. Orcutt, A. Khilo, B. Moss, C. Hozwarth, M. Popovic, H. Li, H. Smit, J. Hoyt, F. Kartner, R. Ram, V. Stojanovic, and K. Asanovic, "Building manycore processor to dram networks with monolithic silicon photonics," in *Proceedings of the Sixteenth Symposium on High Performance Interconnects (HOTI-16)*, august 2008.
- [7] A. Hadke, T. Benavides, S. J. B. Yoo, R. Amirtharajah, and V. Akella, "OCDIMM: Scaling the DRAM memory wall using WDM based optical interconnects," in *Proceedings of 16th IEEE Symposium on High Performance Interconnects (HOTI 2008)*, Palo Alto, CA, august 2008.
- [8] B. Jacob, S. Ng, and D. Wang, *Memory Systems - Cache, DRAM, disk*. Morgan Kaufman Publishers, 2007.
- [9] B. Ganesh, A. Jaleel, D. Wang, and B. Jacob, "Fully-buffered dimm memory architectures: Understanding mechanisms, overheads and scaling," *hpca*, vol. 0, pp. 109–120, 2007.
- [10] D. Wang, B. Ganesh, N. Tuaycharoen, K. Baynes, A. Jaleel, and B. Jacob, "Dramsim: a memory system simulator," *SIGARCH Comput. Archit. News*, vol. 33, no. 4, pp. 100–107, 2005.
- [11] N. Kirman, M. Kirman, R. K. Dokania, J. F. Martinez, A. B. Apsel, M. A. Watkins, and D. H. Albonese, "Leveraging optical technology in future bus-based chip multiprocessors," in *MICRO 39: Proceedings of the 39th Annual IEEE/ACM International Symposium on Microarchitecture*. Washington, DC, USA: IEEE Computer Society, 2006, pp. 492–503.
- [12] Micron, "Technical note: Calculating memory system power for ddr3," 2006.
- [13] H. Cho, P. Kapur, and K. C. Saraswat, "Power comparison between high-speed electrical and optical interconnects for interchip communication," *Journal of Lightwave Technology*, vol. 22, no. 9, p. 2021, 2004. [Online]. Available: <http://jlt.osa.org/abstract.cfm?URI=JLT-22-9-2021>
- [14] D. Vantrease, R. Schreiber, M. Monchiero, M. McLaren, N. P. Jouppi, M. Fiorentino, A. Davis, N. Binkert, R. G. Beausoleil, and J. H. Ahn, "Corona: System implications of emerging nanophotonic technology," in *Proceedings of the 35th International Symposium on Computer Architecture (ISCA-35)*. IEEE Computer Society, 2008, pp. 153–164.
- [15] R. Drost, C. Forrest, B. Guenin, R. Ho, A. V. Krishnamoorthy, D. Cohen, J. E. Cunningham, B. Tourancheau, A. Zingher, A. Chow, G. Lauterbach, and I. Sutherland, "Challenges in building a flat-bandwidth memory hierarchy for a large-scale computer with proximity communication," in *HOTI '05: Proceedings of the 13th Symposium on High Performance Interconnects*. Washington, DC, USA: IEEE Computer Society, 2005, pp. 13–22.
- [16] B. Lemoff, M. Ali, G. Panotopoulos, G. Flower, B. Madhavan, A. Levi, and D. Dolfi, "Maui: enabling fiber-to-the-processor with parallel multiwavelength optical interconnects," *Journal of Lightwave Technology*, vol. 22, no. 9, pp. 2043–2054, Sept. 2004.
- [17] P. Lukowicz, J. Jahns, R. Barbieri, P. Benabes, T. Bierhoff, A. Gauthier, M. Jarczyński, G. Russel, J. Schrage, J. Snowdon, M. Wirz, and G. Troster, "Optoelectronic interconnection technology in the holms system," *IEEE Journal of selected topics in quantum electronics*, 1995.
- [18] Q. Xu, B. Schmidt, S. Pradhan, and M. Lipson, "Micrometre-scale silicon electro-optic modulator," *Nature*, 2005.
- [19] B. A. Small, B. G. Lee, K. Bergman, Q. Xu, , and M. Lipson, "Multiple-wavelength integrated photonic networks based on microring resonator devices," *Journal of Optical Networking*, 2006.
- [20] W. M. J. Green, M. J. Rooks, L. Sekaric, and Y. A. Vlasov, "Ultra-compact, low rf power, 10 gb/s silicon mach-zehnder modulator," *Optical Express*, 2007.