CMTL
Computer Memory Test Labs

HyperCloud HCDIMM Outperforms LRDIMM in 'Big Data' & 'Big Memory' Applications

Whitepaper

March 2012

Document Number: 12107

This paper compares the performance of HyperCloud™ DIMM (HCDIMM™) and Load Reduced DIMM (LRDIMM) in high capacity memory applications used in servers and High-Performance Computing (HPC) platforms. Although targeted for the same types of applications, the architectural implementations approach the problems of memory capacity vs. memory bandwidth and latency in different ways with varying results for bandwidth, capacity, and latency. Comparison data between identically configured next generation Intel® Xeon® Processor E5-2600 systems confirms that HCDIMM significantly outperforms LRDIMM.

Architectural Comparison Between HCDIMM and LRDIMM

As noted earlier, different architectural implementations to solve the electrical loading of the host memory bus and memory sub-system result in superior performance for HCDIMM due to its distributed architecture vs. the single buffer approach of LRDIMM. Key architectural differences are noted below.

HCDIMM Distributed Architecture

The HCDIMM Register re-drives the command, address, and clock signals from the host memory controller to the multiple ranks of DRAM chips and performs the rank multiplication function on the memory module.

The HCDIMM Isolation devices perform the Load Reduction function for the multiple ranks of DRAM chips on the module. A four (4) rank module is thus presented as two virtual ranks (2 vRank) to the host controller memory interface with a significant reduction in the electrical loading seen by the bus and controller.

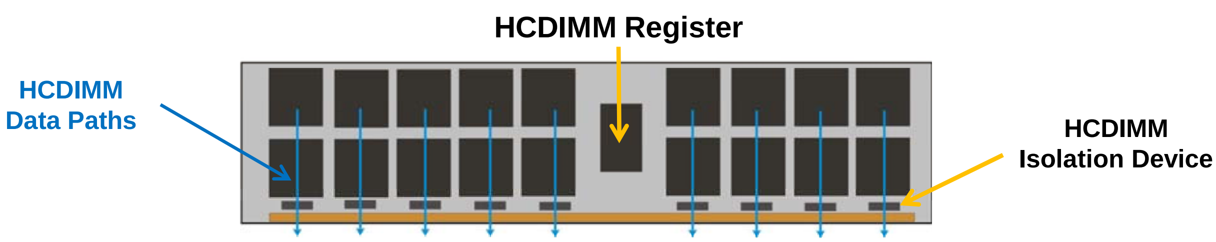

Fig. 1: HCDIMM Distributed Architecture

The HCDIMM distributed architecture is future ready as the industry is adopting the distributed architecture for next generation DDR4 memory modules.

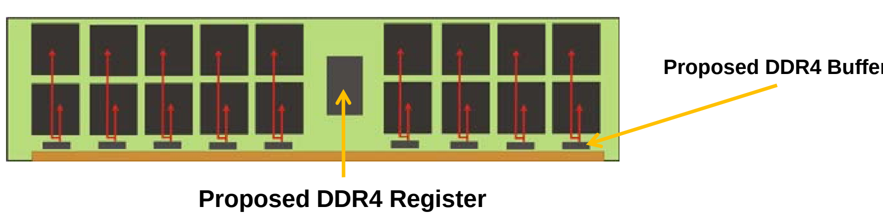

Fig. 2: DDR4 Distributed Architecture

LRDIMM Single Buffer Architecture

Unlike HCDIMM, the LRDIMM approach uses a single device (the memory buffer) to perform both the rank multiplication function and the load isolation function.

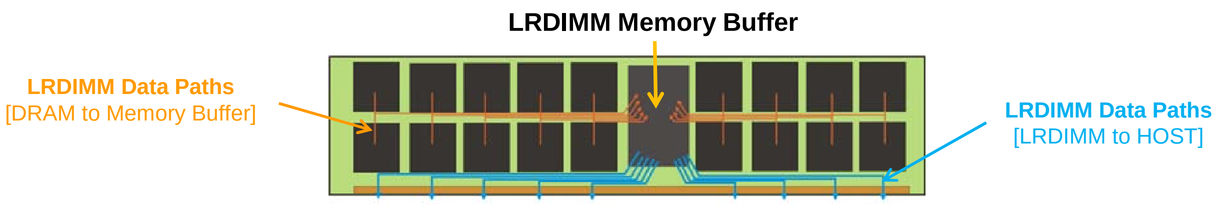

Fig. 3: LRDIMM Singe Buffer Architecture

HCDIMM vs. LRDIMM Architectural Comparison Summary

Architecturally, the HCDIMM implementation will have inherently lower bit-to-bit data skew as well as latency due to its distributed architecture. The data trace routing on the PCB is through multiple Isolation Devices and is more streamlined enabling shorter inline connections between the DRAM ICs through the Isolation Devices to the DIMM connector.

Table 1: HCDIMM vs. LRDIMM Architectural Comparison

	HCDIMM™	LRDIMM	HCDIMM Advantage
Architecture	Separate Register device with distributed Isolation devices (buffers)	Single register device with integrated buffers	Superior HCDIMM performance due to lower latency and skew. Industry adopting Netlist's Distributed Architecture for DDR4
Rank Multiplication	Presents four (4) physical ranks as two virtual ranks (2 vRank) to the host memory controller	Requires BIOS to perform rank reduction	Enables 3DPC at faster memory speeds
Load Reduction	Multiple Isolation devices between DRAM and data bus	Load Isolation integrated into register device requiring fan-in and fan-out of data lines to middle of PCB	HCDIMM has lower data bit-to-bit skew and latency → higher memory bandwidth

The LRDIMM architecture requires data traces to be routed from the DIMM connector to the single memory buffer which are then routed to the DRAM ICs causing multiple data bottlenecks. The first LRDIMM bottleneck is routing the data traces on the PCB itself (similar issue as seen on FBDIMM) and the second bottleneck is routing the data within the memory buffer itself. The LRDIMM data trace length can vary by a factor of 3X to 4X from the shortest trace to the longest trace. In addition to compensating for the data routing length differences within the memory buffer, the memory buffer must add additional latency to compensate for the much higher data skew due to the PCB routing asymmetries.

HCDIMM™ Isolation Devices perform the Load Reduction function with one device placed in-line with each set of DRAM ICs (see Fig 1 for more details) and each device isolates the host memory controller from 4 DRAM electrical loads so that the host sees only 2 electrical loads.

The LRDIMM load isolation is integrated into the memory buffer register device which results in a PCB design where the data traces must fan-in and fan-out from the memory buffer causing additional data bit-to-bit skew and latency.

Table 2: HCDIMM vs. LRDIMM Functional Comparison

HCDIMM™	LRDIMM	ADVANTAGE	System Benefit
Symmetrical data paths on PCB	Asymmetrical data paths on PCB	HCDIMM	Reduced data bit-to-bit skew and latency → higher memory bandwidth
Shorter data paths from DRAM to connector	Longer data paths must be fanned-in to memory buffer and then fanned-out to connector	HCDIMM	Reduced data bit-to-bit skew and latency → higher memory bandwidth
Data paths independent of Register device	Asymmetrical data paths within register device	HCDIMM	Reduced latency → higher memory bandwidth

Architecturally and functionally, the HCDIMM implementation provides a superior solution today and is future-ready for the next generation DDR4 memory implementations.

HCDIMM vs. LRDIMM Performance Comparisons

A dual socket Intel® Xeon® Processor E5-2600 motherboard with two Intel® E5-2650L 1.8GHz processors and 384GB memory configuration (3 DIMM Per Channel, 24 DIMMs total) using 16GB HCDIMM and LRDIMM modules was used as a test bed to evaluate the performance of both memory types under identical conditions by Computer Memory Test Labs, the leading independent compatibility test lab for memory module and motherboard compatibility testing. HCDIMM testing was performed at 1333 MT/s and LRDIMM at 1066 MT/s as systems and motherboards are available to support these configurations as the maximum speeds. SiSoftware® Sandra® Lite benchmarking software was used for memory performance measurements.

HCDIMM vs. LRDIMM Memory Bandwidth Comparisons

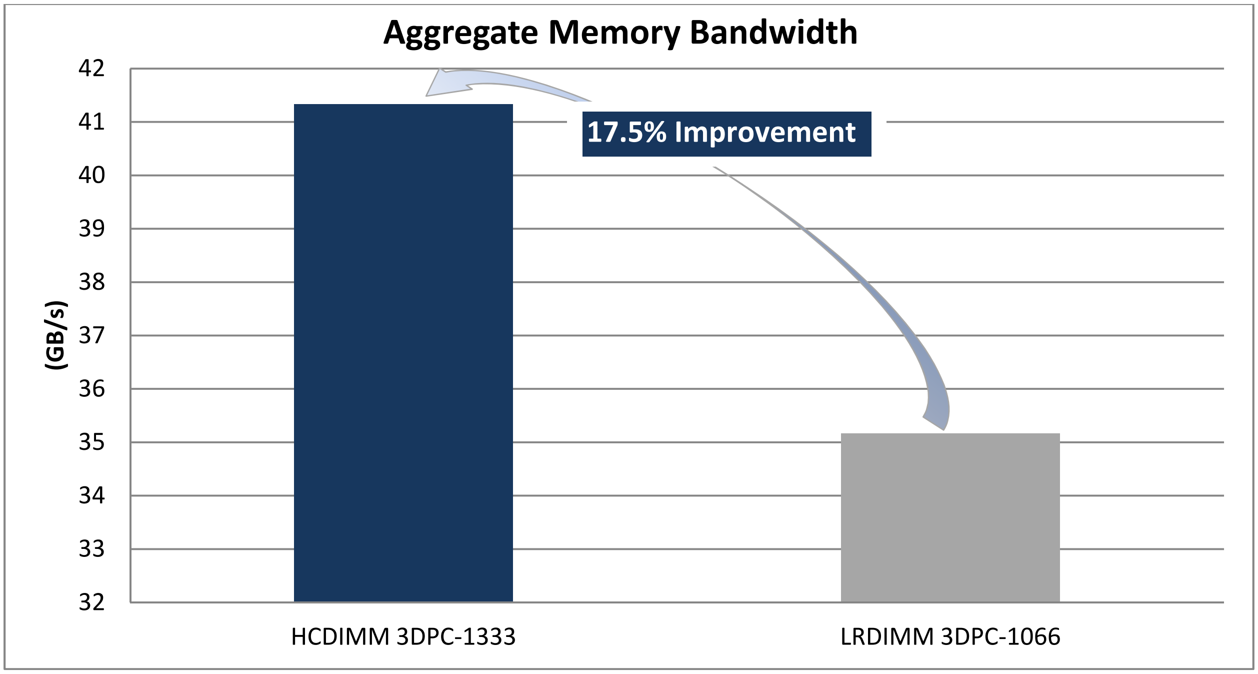

Chart 1: Memory Bandwidth Comparison – (Higher is Better)

HCDIMM vs. LRDIMM Time to Copy Comparisons

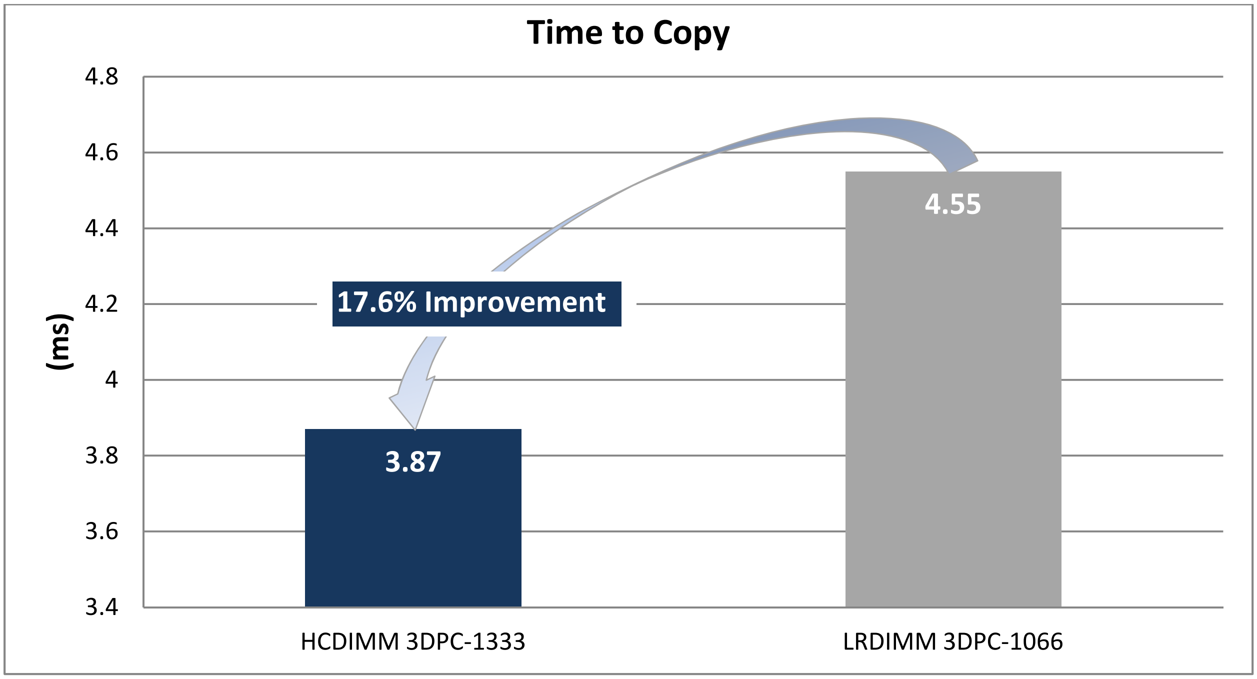

Chart 2: Time to Copy Comparison – (Lower is Better)

HCDIMM vs. LRDIMM Cache Memory Bandwidth Comparison

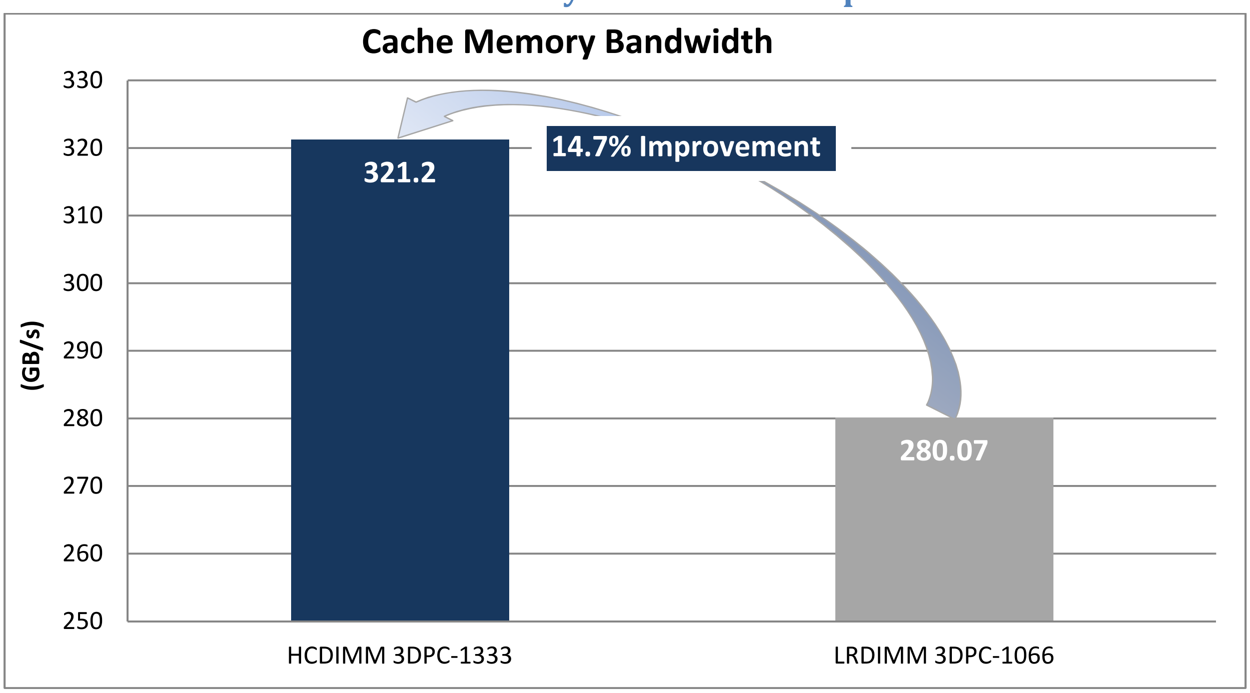

Chart 3. Memory Bandwidth Comparison – (Higher is Better)

Comparison Summary

HyperCloud's Distributed Architecture advantages are highlighted by improved memory bandwidth performance over LRDIMM single memory buffer modules. Since HCDIMM operates at one memory speed grade higher (1333MT/s) in fully configured three DIMMs Per Channel (3DPC) systems, further performance gains are realized vs. LRDIMM which are supported at a maximum of 1066 MT/s operation in 3DPC configurations.

HCDIMM improves 'Aggregate Memory Bandwidth' (a combination of Integer B/F S and Float B/F SSE) over LRDIMM by 17.5%. Similarly, HCDIMM outperforms LRDIMM in the 'Time to Copy' benchmark by 17.6% and the 'Cache Memory Bandwidth' by 14.7%. For Big Data and In Memory DataBase (IMDB) applications, HyperCloud's memory bandwidth performance advantage translates directly into increased productivity and less expenditures in hardware, software, and personnel resources.

The types of applications that can benefit from improved memory performance vary from Finance and Trading to Big Data Analytics, and simulation intensive applications such as EDA (Electronic Design Automation) and FEA (Finite Element Analysis).

Table 4: Simulation Application Benefits / Advantages (EDA, FEA, CFD, CED)

HCDIMM™ Benefits	HCDIMM™ Advantages
❖ Faster simulation runtimes ❖ Larger, more complex simulations ❖ Higher accuracy simulations ❖ Shorter design cycles ❖ Increased number of projects ❖ Reduced simulation re-runs	✓ **Quicker Time to Market** ✓ **Increased Revenue Streams from more project runs** ✓ **Quicker Time to Revenue** ✓ **Reduced OPEX (Operational Expenditures)** ✓ **Reduced CAPEX (Capital Expenditures**

Finance and Trading as well as Big Data Analytics will enjoy similar advantages as simulation intensive applications.

Table 5: Finance and Trading, Big Data Analytics Benefits / Advantages

HCDIMM™ Benefits	HCDIMM™ Advantages
❖ Larger In Memory Databases for more historical data access ❖ Lower latency algorithmic trading ❖ Almost instant data manipulation ❖ Faster concurrent IQ query execution ❖ Reduced data latency ❖ Reduced execution latency	✓ **Higher transaction processing efficiency** ✓ **Reduced per Transaction processing costs** ✓ **Faster Revenue generation** ✓ **Increased Revenue generation** ✓ **More Transaction processing** ✓ **Reduced OPEX** ✓ **Quicker time to Market**

Table 6: Virtualization Benefits / Advantages

HCDIMM™ Benefits	HCDIMM™ Advantages
❖ Increased number of Virtual Machines (VMs) & Virtual Desktop instances (VDI's) ❖ Increased memory per VM and VDI	✓ **Increased productivity** ✓ **Increase hardware utilization / efficiency** ✓ **Reduced Total Cost of Ownership (TCO)**

Summary

HyperCloud Distributed Architecture DDR3 HCDIMMs utilized in high performance server and HPC platforms results in significant memory bandwidth improvements over LRDIMM solutions. Many of the today's key applications will execute faster run times resulting in increased revenue generation, quicker time to market, reduced Operational Expenses (OPEX), and Total Cost of Ownership (TCO) through more efficient hardware, software and personnel utilization.

Technical Information

Server Motherboard:	Super Micro Computer® X9DR6-LN4+
Processors:	Dual Intel® Xeon® CPU E5-2650L 0 @ 1.8GHz
Operating System:	Microsoft® Windows Server 2008®; SP1
HCDIMM Memory Modules:	16GB 2vRank x4; 1333 MT/s
LRDIMM Memory Modules:	16GB 4Rank x4; 1333 MT/s at supported 1066 MT/s in 3DPC
Benchmark Suite:	SiSoftware® Sandra® Lite SP1 version 18.10
Test Facility:	Computer Memory Test Labs – Anaheim, CA