

DDR4 LRDIMMs Let You Have It All

LRDIMMs provide a superior alternative solution for both deeper memory *and* higher data bandwidth

By Douglas Malech and Sameer Kuppahalli, IDT and Ryan Baxter and Eric Caward, Micron Technology, Inc. Endorsed by HP Server Architecture Team

Introduction

Traditionally, LRDIMMs and RDIMMs have provided complementary solutions for data center enterprise servers – LRDIMMs for applications requiring deeper memory and RDIMMs for applications requiring higher data bandwidth. However, with the introduction of 8 gigabit (Gb) DRAMs, this white paper shows that LRDIMMs provide a superior alternative solution for both deeper memory **and** higher data bandwidth. In particular, 32 GB 2RX4 LRDIMMs based on 8 Gb DDR4 DRAMs are shown to transcend 32 GB 2RX4 RDIMMs in both server memory capacity and bandwidth.

The Need for Both Deeper Memories and Higher Data Bandwidth

There are a growing number of Internet applications benefitting from both deeper memories and higher data bandwidth such as in-memory databases (IMDBs). Larger IMDBs mean that more data can reside in high-speed DDR4 DRAM memory, reducing data exchanges with slower storage media during data processing. Reducing data exchanges with slower storage media means the memory-intensive application will run faster. IMDB application examples include data analytics, financial algorithms, gaming and search algorithms. Figure 1 shows some examples.


Figure 1

Application	Problem to Solve	Benefit from Deeper Memory	Benefit from Higher Data Bandwidth
In-memory databases	Speed up database application response time	Keep entire database in memory to reduce Rd/Wrt access to slower storage media	Faster processing means faster results
Social Media	Move text, pictures and video quickly between many users	Caching content for quicker retrieval; e.g. streaming video	Instant response time is desirable
Big Data	Analyze larger and larger data sets in less time	Keep entire dataset in memory to reduce Rd/Wrt access to slower storage media	Faster processing means faster results
Search Algorithms	Faster search algorithms	Keep entire search database in memory to reduce Rd/Wrt access to slower storage	Instant response time desirable
Financial	Speed up financial algorithm response time	Even ms of speed improvement can return millions of dollars in profits	Even ms of speed improvement can return millions of dollars in profits
Virtual Servers	Multiple virtual servers on one physical server	Pairing more memory with more processor cores allows more virtual servers	Speed of virtual servers shouldn't slow down as number of virtual servers goes up
IoT	Real-time multi-sensor continuous transmission	Fast ingress and egress of large sensor data sets with minimal access to slower storage	Stay ahead of real-time data requirements

The LRDIMM Advantage


Prior to 8 Gb DRAMs, 32 GB LRDIMMs were constructed using more expensive DDP¹ packaged 4 Gb DRAMs, as shown in Figure 2. In 32 GB 4RX4 LRDIMMs based on 4 Gb DRAMs, two DRAM data loads from the frontside DDP and two from the backside DDP share the same data bus. These four DRAM data loads are further reduced down to one data buffer load by the LRDIMM's unique data buffering architecture.

Figure 2



Because of this data load reduction technique, when populating three LRDIMMs in a memory channel, only three loads are present. In addition, the nine data buffers are physically located very close to the edge connector, reducing data transmission stub length. Reducing transmission stub lengths and stub counts improve signal integrity. Figure 3 shows a memory controller channel (MCH) configured with three LRDIMMs (3 DPC). Improving signal integrity adds more signal margin, thereby improving server system reliability when populating thousands of memory modules into thousands of servers at a typical data center.

Figure 3



¹ DDP = Dual Die Package. 2 DRAM die in a single DRAM package


Contrary to 32 GB 4RX4 LRDIMMs, 32 GB 4RX4 RDIMMs were not developed because, in the absence of load reducing data buffers, all four DRAM data loads would be visible to the MCH channel, presenting twelve loads in a three RDIMM per MCH configuration (4 DRAM loads x 3 RDIMM). In addition, without data buffers, the signal distance from the DRAMs to the edge connector is increased. An increase in transmission stub lengths and stub counts means poorer signal integrity. This is why RDIMMs based on 4 Gb DRAMs stop at 16 GB 2RX4 memory capacity while LRDIMMs go up to 32 GB 4RX4 memory capacity.

As applications continue to benefit from increased memory capacity, 8 Gb DRAMs enable the RDIMM “sweet spot” to increase from 16 GB memory modules to 32 GB. A 16 GB RDIMM is constructed using 4 Gb DRAMs in a 2RX4 configuration. It follows that a 32 GB RDIMM can be constructed from 8 Gb DRAMs using the same 2RX4 configuration because each DRAM contributes twice as much memory. Likewise, a 32 GB LRDIMM can be reconstructed using 8 Gb DRAMs in a 2RX4 configuration instead of using more expensive DDPs in a 4RX4 configuration. With 8 Gb DRAMs doubling RDIMM memory capacity from 16 GB to 32 GB and simultaneously replacing more expensive DDP’s previously used to construct 32 GB LRDIMMs, which is the better choice for fully populated systems – 32 GB LRDIMMs or 32 GB RDIMMs? Our lab measurements show that 32 GB 2RX4 LRDIMMs have a clear advantage over 32 GB 2RX4 RDIMMs in that you can benefit from the additional memory at a higher bandwidth.

Comparing 32 GB LRDIMMs and 32 GB RDIMMs

A typical enterprise class server can have up to 24 memory modules as shown in Figure 4. A server with 24 memory modules, each with 32 GB of memory, will have 768 Gb of memory (24 x 32 GB).

Figure 4




IDT wanted to see which 32 GB memory module solution provided the superior solution for both total server memory and data bandwidth – LRDIMM or RDIMM. The process by which IDT made this determination was in the following manner:

- 1) Determine the module’s signal integrity on the MCH channel for the fully populated 3 DPC system configuration
- 2) Choose the highest speed possible for LRDIMM and RDIMM with acceptable signal integrity
- 3) Compare bandwidths at these highest speeds to determine whether LRDIMM or RDIMM gives higher memory bandwidth

Data Signal integrity was measured in two places as shown in Figure 5 as $V+$ and $V-$. More positive $V+$ and more negative $V-$ voltage measurements infer better signal integrity. In each case, it is a measure of how much voltage margin is available between the actual Data Eye Signal and the Data Eye Mask. The Data Eye Signal must never dip into the area within the Data Eye Mask for all combinations of data signal patterns, DIMMs, server motherboards and microprocessors. If the Data Eye Signal dips into the Data Eye Mask region, a data value of “1” might be interpreted as a “0” and vice versa.

Figure 5



The four measurements taken for receive and transmit directions show that 3DPC LRDIMM signal integrity at 2400 MT/s operation has better signal integrity than 3DPC RDIMMs at both 2400 MT/s and 2133 MT/s speeds. The measured signal integrity data is shown in Figure 6 with the 3DPC RDIMM measurements in **AMBER** and the 3DPC LRDIMM measurement in **GREEN**. LRDIMM at 2400 MT/s has more positive $V+$ and more negative $V-$, indicating overall better signal integrity.

Figure 6

RDIMM	Receive (DRAM Reads)				Transmit (DRAM Writes)	
	DPC	Speed	V- (mV)	V+ (mV)	V- (mV)	V+ (mV)
	1	2133	-173	164	-156	156
	2	2133	-117	122	-125	125
	3	2133	-66	70	-70	70
	1	2400	-159	159	-133	133
	2	2400	-103	103	-86	94
	3	2400	-51	56	-47	47

LRDIMM	Receive (DRAM Reads)				Transmit (DRAM Writes)	
	DPC	Speed	V- (mV)	V+ (mV)	V- (mV)	V+ (mV)
	1	2133	-173	168	-140	172
	2	2133	-108	98	-156	156
	3	2133	-84	94	-117	117
	1	2400	-164	159	-140	148
	2	2400	-122	117	-140	148
	3	2400	-80	84	-101	109

LRDIMM SI margin improvement at 3DPC

DPC	Speed	Receive (DRAM Reads)		Transmit (DRAM Writes)	
		V- (mV)	V+ (mV)	V- (mV)	V+ (mV)
3	2133	19	23	47	47
3	2400	28	28	55	62

Since 3DPC RDIMMs at 2400 MT/s had much lower voltage margins, IDT assumed that this RDIMM combination of speed and density would be ignored as a possible candidate for server applications. While 3DPC RDIMMs at 2133 MT/s also showed lower voltage margins, IDT chose this 2133 MT/s speed configuration, in absence of the 2400 MT/s option, to compare bandwidths with 32 GB LRDIMMs operating at 2400 MT/s.

IDT used Membw² to compare bandwidths. Membw is a memory bandwidth testing software package that is public domain software. Membw stresses the memory modules with reads and writes across all memory channels. The server configuration used in this benchmarking exercise has two Intel multi-core microprocessors, each with 4 memory channels and 3 DPC for a total of 24 memory modules. The Membw benchmark measurements showed that 3DPC LRDIMM bandwidth at 2400 MT/s is 8% higher than the 3DPC RDIMM bandwidth at 2133 MT/s.

Figure 7

3DPC, 2Rx4, 32GB modules	RDIMM	LRDIMM	% improvement with LR
Speed (MT/s)	2133	2400	13%
Measured Bandwidth (GB/s)	125	135	8%

Conclusion

Evolving enterprise server applications will benefit from both higher bandwidth and more memory module capacity. IDT compared 32 GB 2RX4 LRDIMM and 32 GB 2RX4 RDIMM performance in a 3DPC configuration for both signal integrity and read/write bandwidth. A fully populated server with 24 32 GB LRDIMMs operating at 2400 MT/s showed better signal integrity than the same configuration using 24 32 GB RDIMMs operating at 2133 MT/s. The LRDIMM operating at 2400 MT/s also has an 8% higher bandwidth than RDIMMs operating at 2133 MT/s.

DDR4 LRDIMMs let you achieve higher memory bandwidth than RDIMMs at even mainstream module densities.

² <https://github.com/tomari/membw>

Glossary

16 GB	16 Gigabyte or 16,000,000,000 bytes
2133 MT/s	2,133,000,000 "0/1" pattern transitions travel across the edge connector every second for each of the 64 data bits
2400 MT/s	2,400,000,000 "0/1" pattern transitions travel across the edge connector every second for each of the 64 data bits
2RX4	2RX4 stands for two things. "2R" stands for "dual rank". "X4" stands for DRAMs with 4 DQ data I/O bits
32 GB	32 Gigabyte or 32,000,000,000 bytes
4 Gb	4 Gb or 4,000,000,000 bits
4RX4	4RX4 stands for two things. "4R" stands for "quad rank". "X4" stands for DRAMs with 4 DQ data I/O bits
8 Gb	8 Gb or 8,000,000,000 bits
DB	Bidirectional data buffer. Performs 4:1 load and isolated DRAMs from the motherboard. Registers and re-drives data to and from the DRAMs
DDP	Dual Die Package. 2 DRAM die in a single monolithic chip or DRAM package
DPC	DIMM per channel. In a memory channel, there can be 1, 2, or 3 slots for DIMMs
IMDB	In memory database
IoT	Internet of Things
LRDIMM	Load reduced memory module. Includes an RCD and 9 DB to buffer address, command and data
MCH	Memory channel, usually integrated into the host microprocessor. Some microprocessors have up to four integrated memory channels
mV	millivolts
RCD	Registered clock driver. Buffers and re-drives command and address to the DRAMs
RDIMM	Registered memory module. Includes an RCD to buffer address and command

DISCLAIMER Integrated Device Technology, Inc. (IDT) and its subsidiaries reserve the right to modify the products and/or specifications described herein at any time and at IDT's sole discretion. All information in this document, including descriptions of product features and performance, is subject to change without notice. Performance specifications and the operating parameters of the described products are determined in the independent state and are not guaranteed to perform the same way when installed in customer products. The information contained herein is provided without representation or warranty of any kind, whether express or implied, including, but not limited to, the suitability of IDT's products for any particular purpose, an implied warranty of merchantability, or non-infringement of the intellectual property rights of others. This document is presented only as a guide and does not convey any license under intellectual property rights of IDT or any third parties. IDT's products are not intended for use in life support systems or similar devices where the failure or malfunction of an IDT product can be reasonably expected to significantly affect the health or safety of users. Anyone using an IDT product in such a manner does so at their own risk, absent an express, written agreement by IDT.

IDT and the IDT logo are registered trademarks of Integrated Device Technology, Inc., in the United States and other countries. All other trademarks are the property of their respective owners. © 2016. Integrated Device Technology, Inc. All Rights Reserved.

©2016 Micron Technology, Inc. All rights reserved. Information, products, and/or specifications are subject to change without notice. All information is provided on an "AS IS" basis without warranties of any kind. Drawings may not be to scale. Micron, the Micron logo, and all other Micron trademarks are the property of Micron Technology, Inc. All other trademarks are the property of their respective owners.