

TMS320TCI6482 EDMA3 Performance

Brighton Feng, Rafeeq Salman

ABSTRACT

The enhanced DMA v3 (EDMA3) controller of the TMS320TCI6482 device is a highly efficient data transfer engine, capable of maintaining transfers at up to 21 GB/sec at a 1 GHz CPU clock frequency. This document details measured bandwidth achieved under various operating conditions. For more information on ideal transfer bandwidth and scheduling transfers, please consult *TMS320TCI648x DSP Enhanced DMA (EDMA3) Controller User's Guide* ([SPRU727](#)).

Contents

1	Introduction	2
2	Transfer Overhead	2
3	Max Bandwidth Between Different Memories	2
4	Bandwidth vs. Transfer Flexibility	4
5	EDMA3 vs. IDMA	7
6	EDMA3 vs. CPU for Data Transfer.....	7
7	Summary	11
8	References	11

List of Figures

1	Burst Size Utilization	4
2	Linear 2D Transfer	5
3	Index Effect on EDMA Bandwidth	6
4	CPU Reads From DDR2 Using LDW Instructions	8
5	CPU Reads From DDR2 Using LDDW Instructions	9
6	CPU Writes to DDR2 Using STW Instructions	10
7	CPU Writes to DDR2 Using STDW Instructions	11

List of Tables

1	Transfer Overhead	2
2	Max Bandwidth Between Different Memories.....	3
3	Difference Between TCs.....	3
4	Throughput Comparison Between TCs	3
5	Throughput Comparison Between EDMA3 and IDMA.....	7
6	Throughput Comparison Between EDMA and CPU	7

1 Introduction

The enhanced DMA (EDMA3) controller of the TMS320C64x+ devices includes four parallel highly efficient data transfer engines, each of them capable of handling up to 16 bytes per EDMA3 cycle, resulting in about 21 GB per second of total data throughput at a CPU rate of 1 GHz (the EDMA3 frequency being CPU frequency divided by three). The EDMA3 performs all data movement between the on-chip level-one (L1) memory, level-two (L2) memory, external memory, and the device peripherals. These data transfers include CPU-initiated and event-triggered transfers. The EDMA3 architecture has many features designed to facilitate simultaneous multiple high-speed data transfers. With a working knowledge of this architecture and the way in which data transfers interact and are performed, it is possible to create an efficient system and maximize the bandwidth utilization of the EDMA3.

This document gives designers a basis for estimating system performance. Most of the tests operate under best-case situations to estimate maximum throughput that can be obtained. The transfers described in this document serve as a sample of interesting or typical performance conditions. The performance of other configurations is neglected in the interest of brevity and can generally be inferred from the examples given. This document provides a good source of actual performance values to estimate the performance of a system. Most of following performance data is examined on the 1 GHz TCI6482 verification and debugging board (VDB).

2 Transfer Overhead

Initial Latency is defined as the time interval between the triggering of EDMA event and the instance at which the data transfer actually begins. Since initial latency is hard to measure, we measured transfer overhead instead; it is defined as the sum of the latency, epilog and prolog of transfer pipeline. The values vary based on the type of source/destination peripheral and readiness of the source/destination ports and peripherals. The transfer time for smallest element (1 byte) is considered as the transfer overhead. [Table 1](#) shows the cycles consumed for smallest transfer between different ports.

Table 1. Transfer Overhead

Source\Destination	L1	L2	DDR2	McBSP
L1	129	129	145	145
L2	135	135	153	145
DDR2	177	177	193	257
McBSP	190	190	200	

As conclusion, transfer overhead is a big concern for short transfers and need to be included when scheduling EDMA3 traffic in a system. Single-element transfer performance will be latency-dominated.

3 Max Bandwidth Between Different Memories

The transfer bandwidth is found by taking total bytes transferred and dividing it by the time it is used. In all the cases, transfer bandwidth is limited by the lesser of the source or destination port bandwidth. The maximum speed the transfer could achieve is equal to the bandwidth of the limiting port. The four transfer engines of EDMA3 are not exactly the same. In some cases, the difference can impact the throughput.

The max bandwidths that can be achieved for different memories are examined on TCI6482 VDB. On the VDB, A 32-bit width, 250 MHz DDR2 is connected to the DSP. The actual DDR2 data rate is 500 M since the DDR2 can access data on dual clock edges. The internal L2 SRAM runs at ½ CPU clock (500 MHz). The internal L1 SRAM runs at CPU clock (1000 MHz). EDMA3 bus runs at 1/3 CPU clock (~333 MHz).

Table 2 shows the max bandwidth between L1D, L1P, L2 and DDR2.

Table 2. Max Bandwidth Between Different Memories

	Actual Bandwidth (MB/s)	Theoretical Bandwidth (MB/s)	Memo
L1P ↔ L1D	3799	5333	Theoretical EDMA3 bandwidth is: (128/8)*(1000/3)= 5333 (MB/s)
L1P ↔ L1P	3507	5333	
L1D ↔ L1D	3799	5333	
L1P ↔ L2	4146	5333	
L1D ↔ L2	4146	5333	
L2 ↔ L2	3375	5333	
L2 ↔ DDR2	1989	2000	Theoretical DDR2 bandwidth is: (32/8)*500= 2000 (MB/s)
DDR2 ↔ DDR2	819	1000	

The theoretical bandwidth doesn't consider any overhead. The key factor which reduces the internal memory bandwidth performance is the snoop operation which keeps data coherency between L1 and L2 memory.

For linear internal memory block transfer, IDMA can achieve higher throughput than EDMA3, this is discussed further in the following section.

3.1 Difference Between Four Transfer Engines

EDMA3 includes four (4) TC (transfer controller). If 4 TCs runs simultaneously, the overall throughput for EDMA3 is $4 * (128/8 * 1000/3) = 21333$ MB/s.

The 4 transfer engines are not exactly the same. Table 3 is a summary of the difference.

Table 3. Difference Between TCs

Name	TC0	TC1	TC2	TC3
FIFO Size	128 bytes	128 bytes	256 bytes	256 bytes
Bus Width	128 bits	128 bits	128 bits	128 bits
Destination FIFO Entries	2 entries	4 entries	4 entries	4 entries
Default Burst Size	64 bytes	64 bytes	64 bytes	64 bytes

For more information about the differences, please refer to *TMS320TCI648x DSP Enhanced DMA (EDMA3) Controller User's Guide* ([SPRU727](#)).

Table 4 compares the throughput test result for different TCs.

Table 4. Throughput Comparison Between TCs

Name	TC0	TC1	TC2	TC3
L2 → L1D	2119	2119	3798	3798
L2 ↔ L2	1772	1772	3375	3375
L2 → DDR2	1989	1988	1989	1989
DDR2 → L2	1023	1023	1893	1893
DDR2 ↔ DDR2	815	815	819	819

In conclusion, TC3 and TC2 should have the best performance, followed by TC1, and finally TC0. If no specific note, all performance data in this application report is measured on TC3.

4 Bandwidth vs. Transfer Flexibility

EDMA3 channel parameters allow many different transfer configurations. Most typical transfers burst properly, and memory bandwidth is fully utilized. However, in some less common configurations, transfers are unable to burst, reducing performance. To properly design a system, it is important to know which configurations offer the best performance for high speed operations, and which must trade throughput for flexibility. These considerations are especially important for long memory transfers.

4.1 First Dimension Size (ACNT) Considerations, Burst width

To make full utilization of bandwidth in the transfer engine, it is important to fully utilize the bus width available and allow for data bursting.

ACNT size should be a multiple of 16-byte to fully utilize 128-bit bus width. ACNT should be a multiple of 64-byte to fully utilize 64-byte default burst width. ACNT should be a multiple of 256-byte to fully utilize 256-byte FIFO.

Figure 1 shows performance data from a C64x+ CPU running at 1 GHz, transferring 1~1024 bytes from L2 to DDR2 using an EDMA3 channel.

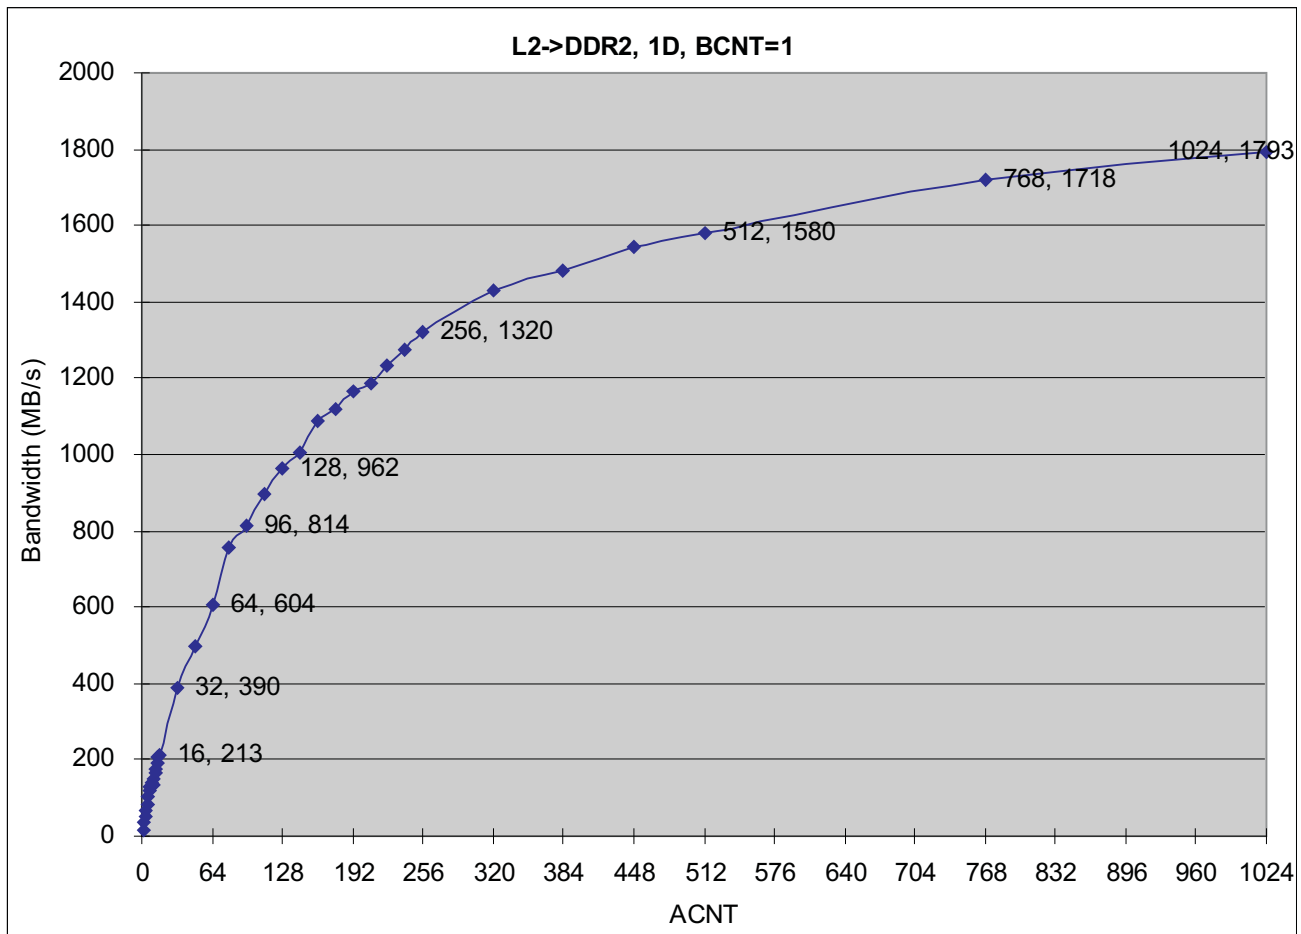


Figure 1. Burst Size Utilization

In conclusion, more bandwidth can be achieved by utilizing larger ACNT.

4.2 Two Dimension Considerations, Transfer Optimization

If 2D transfer (AB_Sync) is linear (BIDX = ACNT), the 2D transfer is optimized as a 1D transfer.

Various ACNT and BCNT combinations were investigated; however, the overall transfer size (ACNT * BCNT) was proved to have more bearing than the particular combination settings. Figure 2 is a linear 2D transfer test result. It shows, no matter what the BCNT is, the bandwidths are similar as long as ACNT * BCNT are the same.

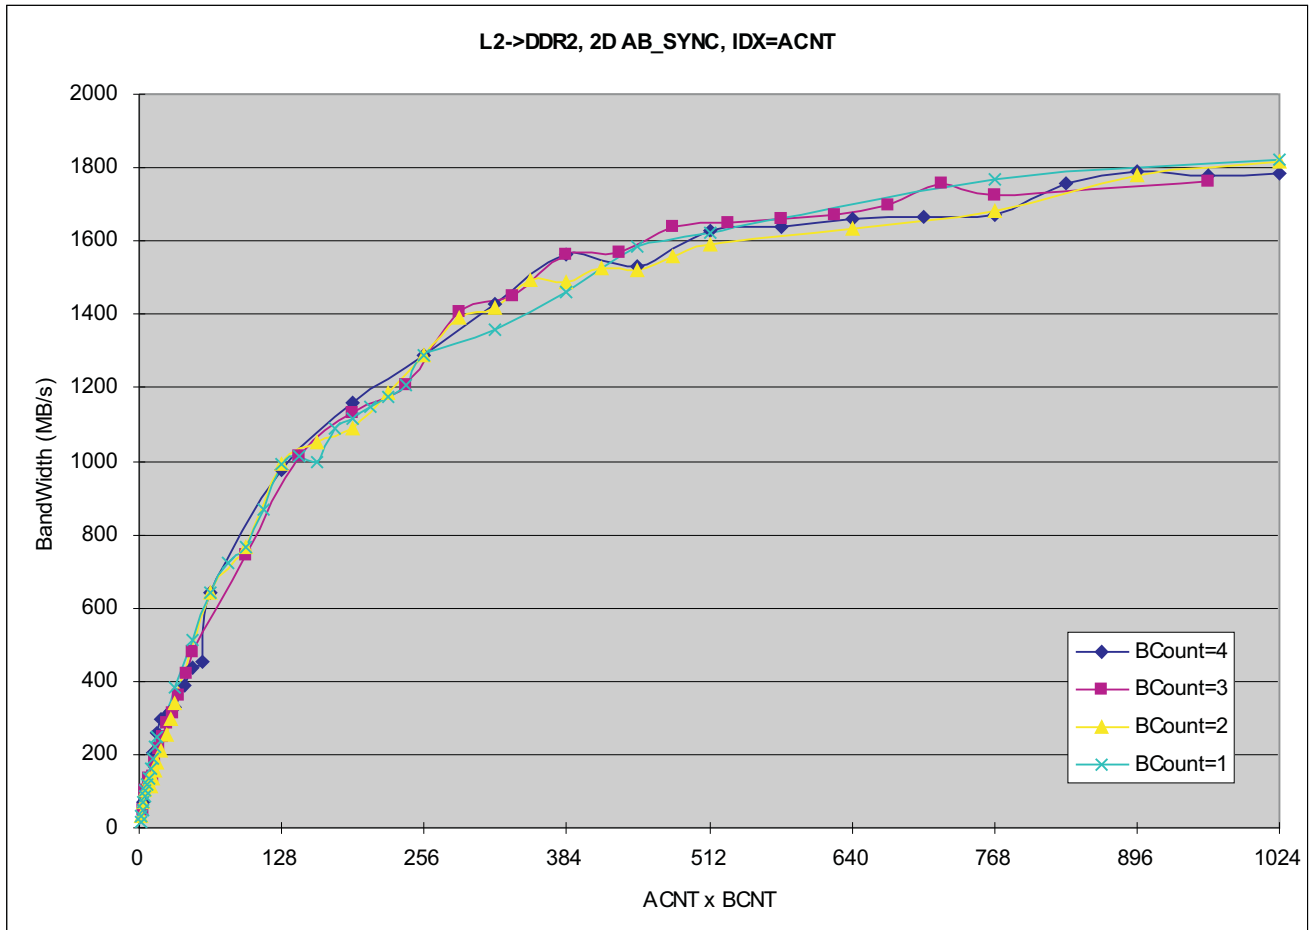


Figure 2. Linear 2D Transfer

If 2D transfer is not linear, the bandwidth utilization is only determined by the ACNT as shown in Figure 1.

4.3 Index Consideration

Index dramatically affects the EDMA throughput. Linear transfer (index = ACNT) fully utilizes bandwidth. Fixed index (index = 0) can utilize almost the same bandwidth as linear transfer. Other index modes lower the EDMA performance. Odd index has the worst performance. If index is a power of 2, and it is larger than 8, then the performance degradation is very small.

Figure 3 shows the index effect on EDMA throughput, transferring 1024 rows (BCNT = 1024) of 2D data from L2 to DDR2, with different index.

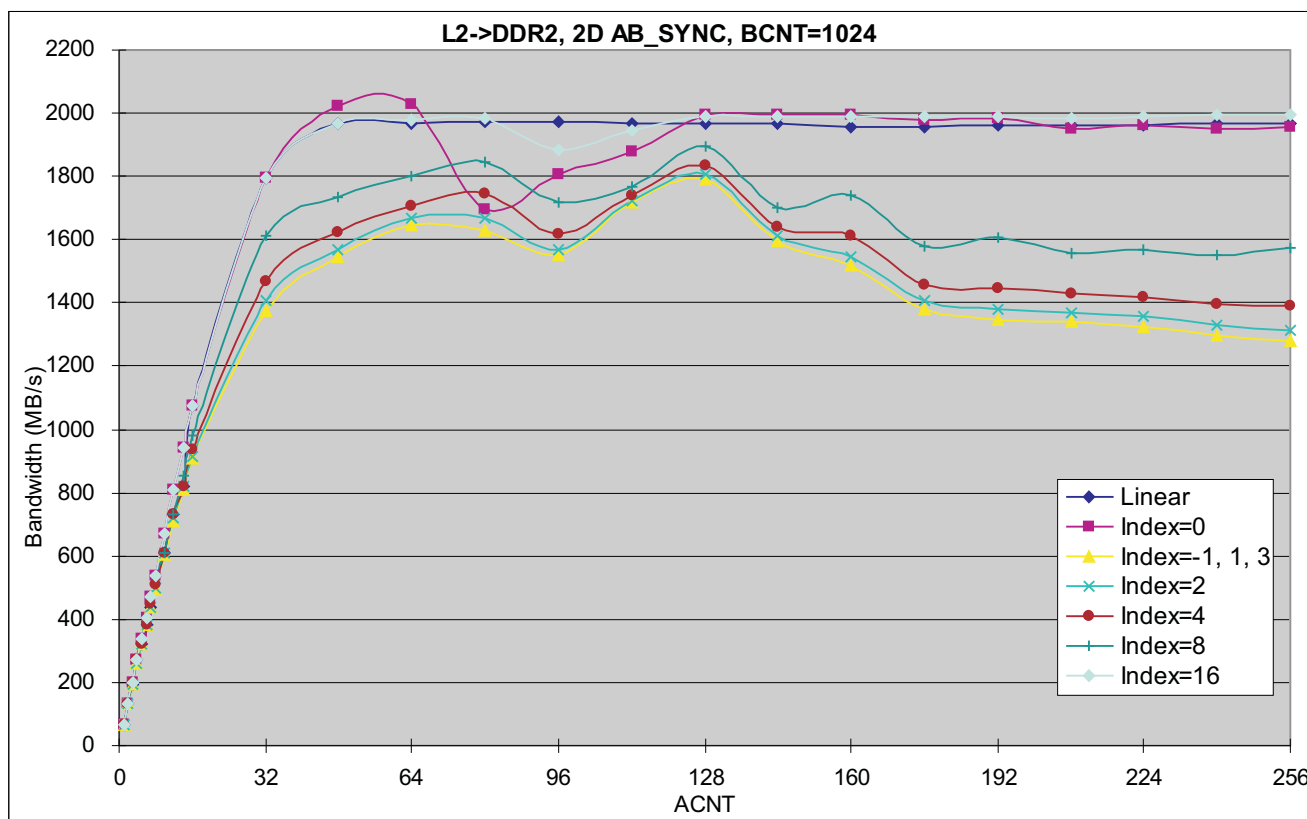


Figure 3. Index Effect on EDMA Bandwidth

4.4 Address Alignment

Address alignment may slightly impact the performance. The default burst size of EDMA3 is 64bytes. If the transfer is across the 64 byte boundary, then the EDMA3TC breaks the ACNT array into 64-byte bursts to the source/destination addresses. So, if the source or destination address is not aligned to the 64 byte boundary and the transfer is across the 64 byte boundary, then extra bursts are generated to handle the unaligned head and tail data.

For big transfer this overhead may be ignored. All data presented in this document are based on the address aligned transfers.

5 EDMA3 vs. IDMA

IDMA is an internal data transfer engine in C64x+ Megamodule. It can only transfer data between L1, L2 and peripheral configuration port. The transfer of IDMA must be linear.

IDMA bus is 256 bits running at ½ CPU clock. IDMA is much faster than EDMA3 for internal memory data transfer. It's latency is much smaller than EDMA3, which means it can be more quickly kicked off.

Test data shows the total transfer overhead of IDMA is about 70 cycles.

Table 5 shows the bandwidth comparison of IDMA and EDMA between different internal memories.

Table 5. Throughput Comparison Between EDMA3 and IDMA

	IDMA Bandwidth (MB/s)	EDMA3 Bandwidth (MB/s)
L2 → L1D	7693.3	3798
L1D → L2	10054.7	4146
L2 → L1P	7693.3	3798
L2 ↔ L2	3526.7	3375
L1D ↔ L1D	3933.8	3799

In conclusion, for the linear data transfer between L1 and L2, IDMA is preferred.

6 EDMA3 vs. CPU for Data Transfer

Generally speaking, using the CPU to access external data is a bad use of resources and should be avoided. Instead, the EDMA3 should be given the task of transferring data so the CPU is free to perform actual computation. EDMA is better for big block of memory transfer.

Table 6 shows the test result of EDMA and CPU for linear 256 KB memory block copy.

Table 6. Throughput Comparison Between EDMA and CPU

Max Bandwidth(MB/s)	EDMA	CPU
L2 → L2 (Unallocated in L1D)	3375	2460
L2 → DDR2 (Cacheable)	1989	575
L2 → DDR2 (NonCacheable)	1989	444
DDR2 → L2 (Cacheable)	1893	967
DDR2 → L2 (NonCacheable)	1893	151
DDR2 → DDR2 (Cacheable)	819	348
DDR2 → DDR2 (NonCacheable)	819	113

As mentioned in the previous section, EDMA transfer overhead is in terms of several hundreds of cycles. So, for smaller transfers, there should be a trade off between EDMA and CPU.

CPU access highly depends on the cache. When the CPU accesses external memory spaces, a transfer request through EMC may be generated (depending on whether the data is cached). The TR is for one of the following: a single element - if the memory space is non-cacheable, an L1 cache line - if the memory space is cacheable and the L2 cache is disabled, or an L2 cache line - if the memory space is cacheable and L2 cache is enabled. No transfer request is generated in the case of an L1 or L2 cache hit.

An external memory can be cached by L1 cache, L2 cache, or neither. If the appropriate MAR bit for a memory space is not set, it is not cacheable. If the MAR bit is set and L2 cache size is zero (all L2 is defined as SRAM), the external memory space is cached by L1. If the MAR bit is set and L2 cache size is greater than 0, the external memory space is cached by L2. Remember that L2 is always cached by L1.

The address increment (or memory stride) controls cache utilization. Contiguous LDW or STW accesses (a stride of one word or 4 bytes) utilize cache memory to the fullest. A memory stride of 64 bytes or more causes every access to miss in the L1 cache because the L1 line size is 64 bytes. A memory stride of 128 bytes causes every access to miss in L2 because the L2 line size is 128 bytes. Access patterns exploring each of these conditions are covered in the next two sections.

The following sections examine the access time for various CPU accesses to external memory.

6.1 CPU Reads (Load Instructions)

For reads, the CPU stalls, waiting for the return data. The length of the stall is equal to the sum of the transfer latency, transfer duration, data return time, and some small cache request overhead.

Figure 4 shows data collected from 1 GHz TCI6482 VDB. The time required for 8192 LDW (load word) instructions was measured, and the average time for each instruction is reported.

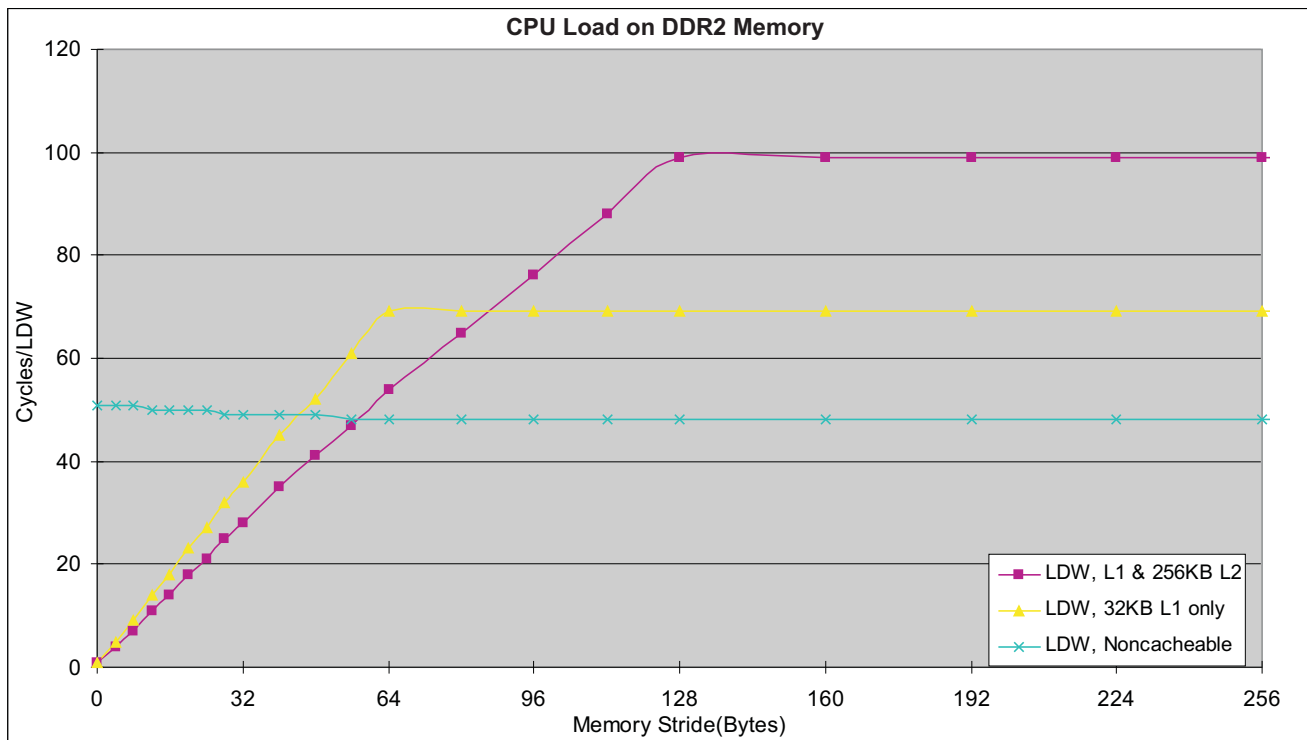


Figure 4. CPU Reads From DDR2 Using LDW Instructions

Figure 5 shows data collected from 1 GHz TCI6482 VDB. The time required for 8192 LDDW (load double word) instructions was measured, and the average time for each instruction is reported.




Figure 5. CPU Reads From DDR2 Using LDDW Instructions

6.2 CPU Writes (Store Instructions)

For writes, a stall may occur if more than four writes are submitted in quick succession. The length of the stall is equal to the sum of the transfer latency, transfer duration, and some small cache overhead. No data return delay is necessary.

Figure 6 shows data collected from a 1 GHz C64x+ writing to 32-bit, 250 MHz external DDR2 memory. The time required for 8192 STW (store word) instructions was measured, and the average time for each instruction is reported.




Figure 6. CPU Writes to DDR2 Using STW Instructions

Figure 7 shows data collected from a 1 GHz C64x+ writing to 32-bit, 250 MHz external DDR2 memory. The time required for 8192 STDW (store double word) instructions was measured, and the average time for each instruction is reported.

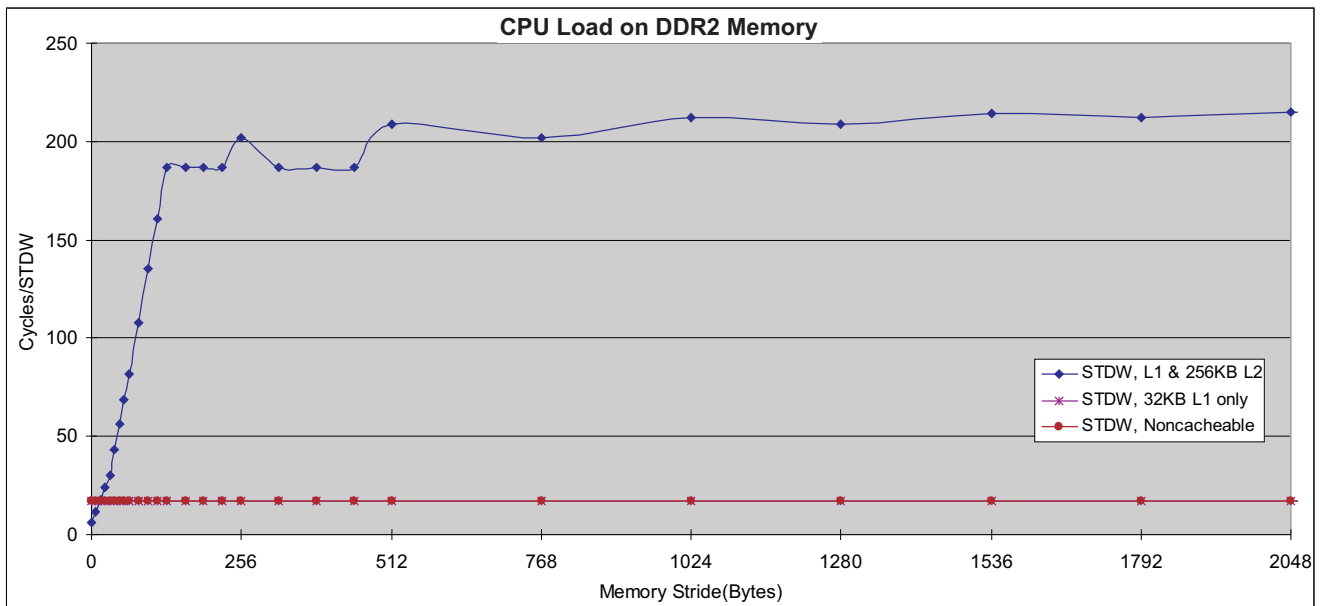


Figure 7. CPU Writes to DDR2 Using STDW Instructions

Since the L1 is a write-through cache, it has no effect on STW/STDW instructions. Therefore, "32KB L1 cache only" case and "non-cacheable" case is overlapped.

Also note, if memory stride is a multiple of 512, only some of the cache lines are utilized, that is why the performance curve is not linear for "L1 & 256 KB L2" case.

7 Summary

This document presented data from EDMA3 throughput trials from a variety of tests and provided insight into the cause of certain performance characteristics. The numbers presented represent measured delays under the circumstances described. These should give designers a good starting point for estimating system performance. Actual performance is application-specific and can vary from these numbers, refer to *TMS320TCI648x DSP Enhanced DMA (EDMA3) Controller User's Guide (SPRU727)* for more information on creating efficient data transfers in a system.

8 References

- *TMS320TCI648x DSP Enhanced DMA (EDMA3) Controller User's Guide (SPRU727)*
- *TMS320C64x+ DSP Megamodule Peripherals Reference Guide (SPRU871)*
- *TMS320C64x EDMA Performance Data Application Report (SPRAA02)*

IMPORTANT NOTICE

Texas Instruments Incorporated and its subsidiaries (TI) reserve the right to make corrections, modifications, enhancements, improvements, and other changes to its products and services at any time and to discontinue any product or service without notice. Customers should obtain the latest relevant information before placing orders and should verify that such information is current and complete. All products are sold subject to TI's terms and conditions of sale supplied at the time of order acknowledgment.

TI warrants performance of its hardware products to the specifications applicable at the time of sale in accordance with TI's standard warranty. Testing and other quality control techniques are used to the extent TI deems necessary to support this warranty. Except where mandated by government requirements, testing of all parameters of each product is not necessarily performed.

TI assumes no liability for applications assistance or customer product design. Customers are responsible for their products and applications using TI components. To minimize the risks associated with customer products and applications, customers should provide adequate design and operating safeguards.

TI does not warrant or represent that any license, either express or implied, is granted under any TI patent right, copyright, mask work right, or other TI intellectual property right relating to any combination, machine, or process in which TI products or services are used. Information published by TI regarding third-party products or services does not constitute a license from TI to use such products or services or a warranty or endorsement thereof. Use of such information may require a license from a third party under the patents or other intellectual property of the third party, or a license from TI under the patents or other intellectual property of TI.

Reproduction of information in TI data books or data sheets is permissible only if reproduction is without alteration and is accompanied by all associated warranties, conditions, limitations, and notices. Reproduction of this information with alteration is an unfair and deceptive business practice. TI is not responsible or liable for such altered documentation.

Resale of TI products or services with statements different from or beyond the parameters stated by TI for that product or service voids all express and any implied warranties for the associated TI product or service and is an unfair and deceptive business practice. TI is not responsible or liable for any such statements.

Following are URLs where you can obtain information on other Texas Instruments products and application solutions:

Products

Amplifiers	amplifier.ti.com
Data Converters	dataconverter.ti.com
DSP	dsp.ti.com
Interface	interface.ti.com
Logic	logic.ti.com
Power Mgmt	power.ti.com
Microcontrollers	microcontroller.ti.com
Low Power Wireless	www.ti.com/lpw

Applications

Audio	www.ti.com/audio
Automotive	www.ti.com/automotive
Broadband	www.ti.com/broadband
Digital Control	www.ti.com/digitalcontrol
Military	www.ti.com/military
Optical Networking	www.ti.com/opticalnetwork
Security	www.ti.com/security
Telephony	www.ti.com/telephony
Video & Imaging	www.ti.com/video
Wireless	www.ti.com/wireless

Mailing Address: Texas Instruments
Post Office Box 655303 Dallas, Texas 75265